Image Compression Refresher – JPEG 2000 and 3D

David Clunie
CTO, RadPharm

Why compress ?

- Reading speed
  - Retrieval of current & priors (not pre-fetched)
  - Bandwidth & hardware cost
- Archival speed and size
  - Media, power, HVAC, physical space cost
  - Offsite replication bandwidth cost
- Tele-radiology speed & throughput
  - Bandwidth cost
  - Impossible otherwise

Lossless Compression

- Same as “reversible” compression
- Does not compromise image integrity
- Takes advantage of “redundancy” in image
- Why not always use it ?
  - Takes time to encode & decode
  - Only modest saving in terms of size
  - Transmission - both ends have to support it
  - Archival - future software needs to be able to read it
- Use only STANDARD schemes (DICOM)
- Avoid proprietary compression like the plague

Lossy Compression

- Same as “irreversible” compression
- ***<u>Always</u>*** compromises image integrity, by definition
- “Visually lossless” compression
  - still lossy, by definition
  - visually indiscernible loss may still affect performance
- Dubious for long term archival, if not interpreted compressed
  - legal preference is to archive what one interpreted
- Dubious for primary interpretation
  - can/should one throw away information without using it ?
- Dubious for portable media
  - may not be sufficient for referring physician
  - may not be suitable as priors
- Currently forbidden in some applications (e.g., mammo)

Why lossy compress ?

- Lower costs
  - Less disk space in archive
  - Cheaper bandwidth (local or remote)
  - Lowering quality of care to save money ?
- Better argument - enable use-cases
  - Tele-radiology from remote regions with slow connections
  - For applications in which rapid turn-around of interpretation is required

How much loss is OK ?

- It depends
  - Modality
    - e.g., X-rays compress more than CTs
  - Task
    - e.g., tube placement vs. interstitial disease
  - Scheme
    - JPEG (DCT), JPEG 2000 (Wavelet), …
- How much degradation is acceptable ?
  - “non-inferiority” to without compression
  - reduction in performance of specified amount
    - reduction in specificity, sensitivity, AROC, …
  - no visually discernible difference

No rules to guide you ...

- ACR guidelines leave it to radiologist’s discretion
- Vendors offer it without addressing its safety
- FDA does not prohibit it
  - except for mammo primary interpretation

Interpreting the literature

- Perceived quality studies
  - task is to discern visual difference
  - easy, generalizable, but …
  - … of little value in decision making

- Observer-performance studies
  - readers complete a detection or characterization task (find lesions, decide if malignant)
  - hard, expensive, difficult to generalize

Statistical power of studies

- “No difference was found …”
- Was this because
  - there is no difference (of greater than a pre-specified amount) ?
  - there were insufficient readers or cases to detect a difference ?
- Review compression studies with respect to estimates of power !
- Statistical vs. clinically meaningful difference

Few answers here

- Few unequivocal conclusions from the current literature with respect to whether or not any lossy compression at any ratio with any scheme is safe to use
- Limited studies suggest promise though
- Watch for publication of the Canadian study
  - Koff et al
  - 7 regions,5 modalities,3 ratios,80 readers,70 images each
  - no difference at 8-15:1 small and 20-30:1 large images
  - JPEG or JPEG 2000
  - may still not satisfy you
  - see it presented at SSG15-03 Tuesday 10:50am S404CD

Some practical examples

- Mammography and CT
- Lossless performance
  - Is it good enough for most purposes
  - Indicates how much compression is possible
- Lossy compression
  - Examine wavelets, specifically JPEG 2000
  - Some comparison with JPEG (DCT)
- Possibility of 3D compression
  - Exploit redundancy in inter-slice dimension
  - Made possible by JPEG 2000 and adopted in DICOM

Compression Ratio confusion

- Ratio of what relative to what ?
- Number of bits on disk (16), or
- Number of meaningful bits (e.g., 12)

to

- Number of compressed “bits per pixel”

- E.g. 1 bpp - express as 16:1 or 12:1 ?

Lossless mammo compression

- 20 pairs (40 images)
  - Of For Processing and For Presentation
- Three vendors
  - 4 pairs Lorad (1 patient, 4 views)
  - 4 pairs Fischer (1 patient, 4 views)
  - 12 pairs GE (3 patients, 4 views each)
- Images *not* first cropped to breast size

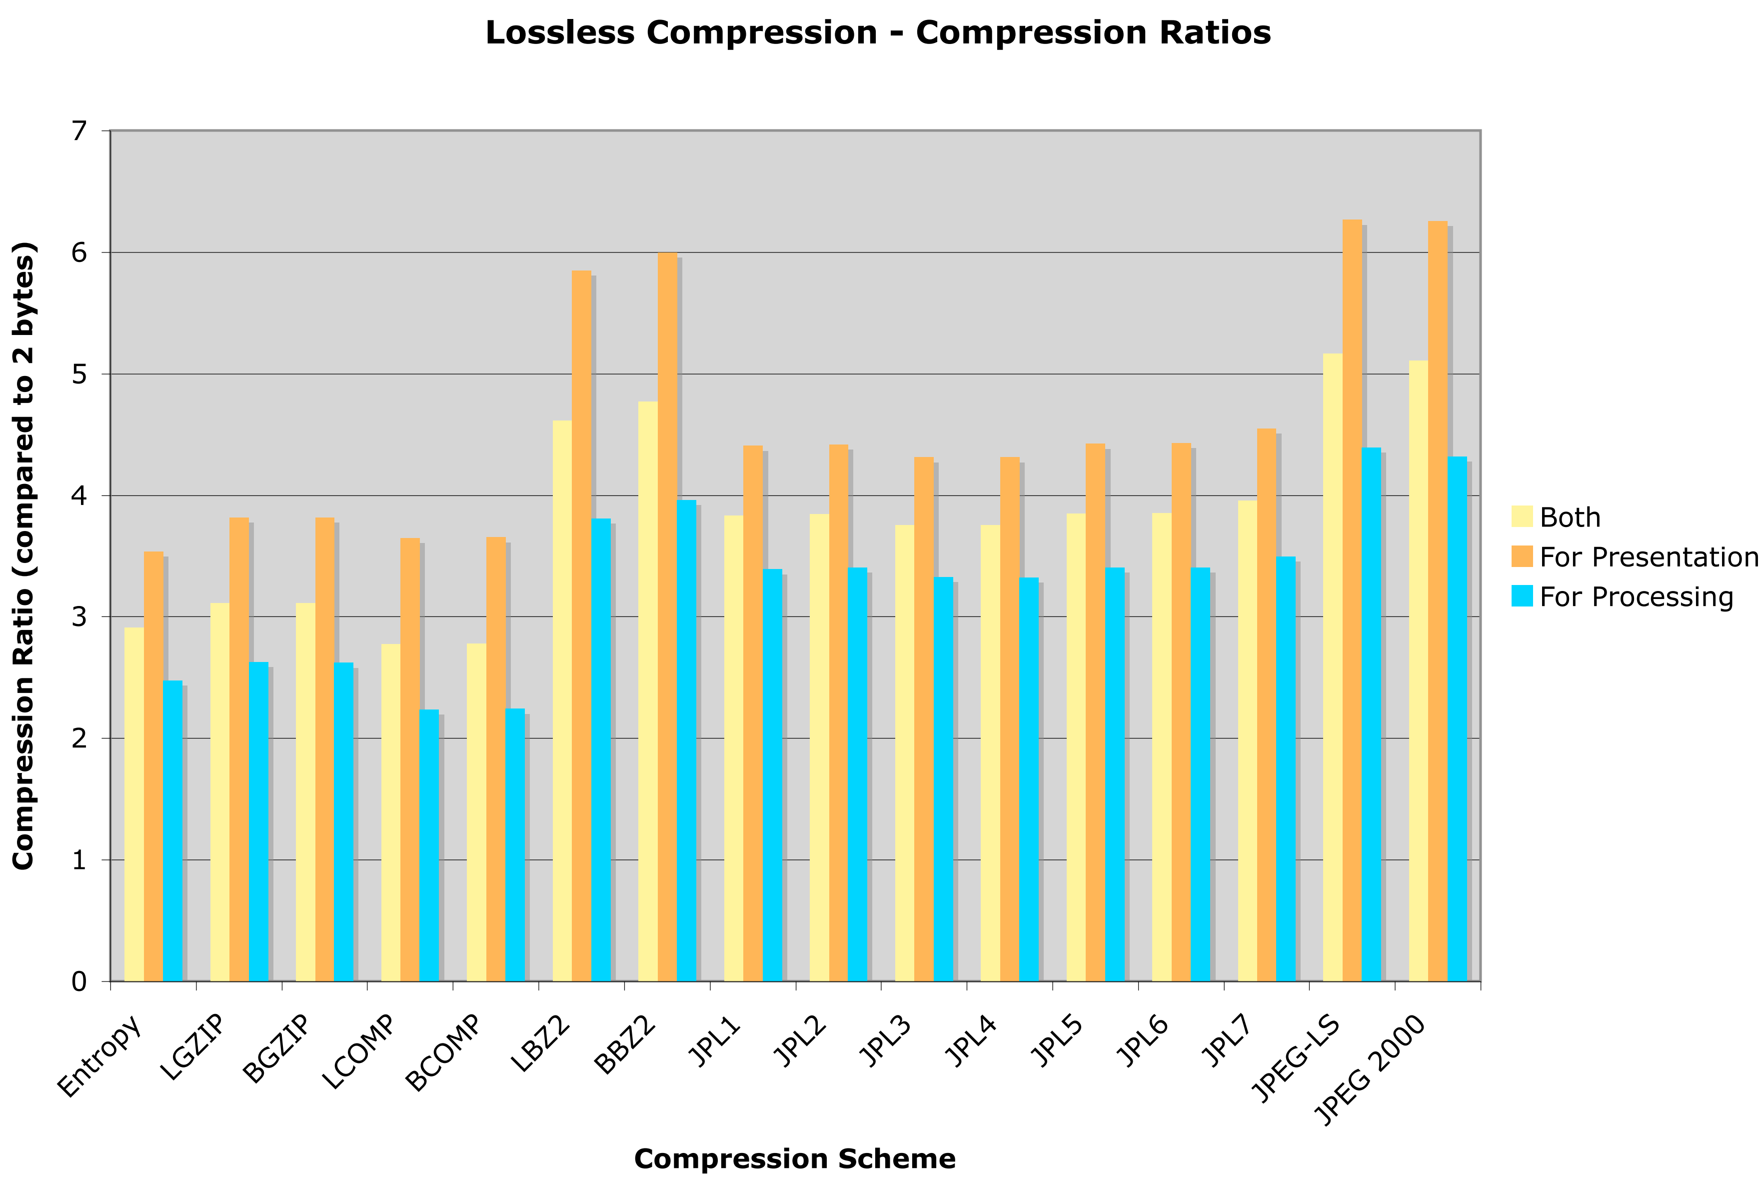

Lossless Compression - Compression Ratios
Compression Ratio (compared to 2 bytes)
0
1
2
3
4
5
6
7
Entropy
LGZIP
BGZIP
LCOMP
BCOMP
LBZ2
BBZ2
JPL1
JPL2
JPL3
JPL4
JPL5
JPL6
JPL7
JPEG-LS
JPEG 2000
Compression Scheme
Both
For Presentation
For Processing

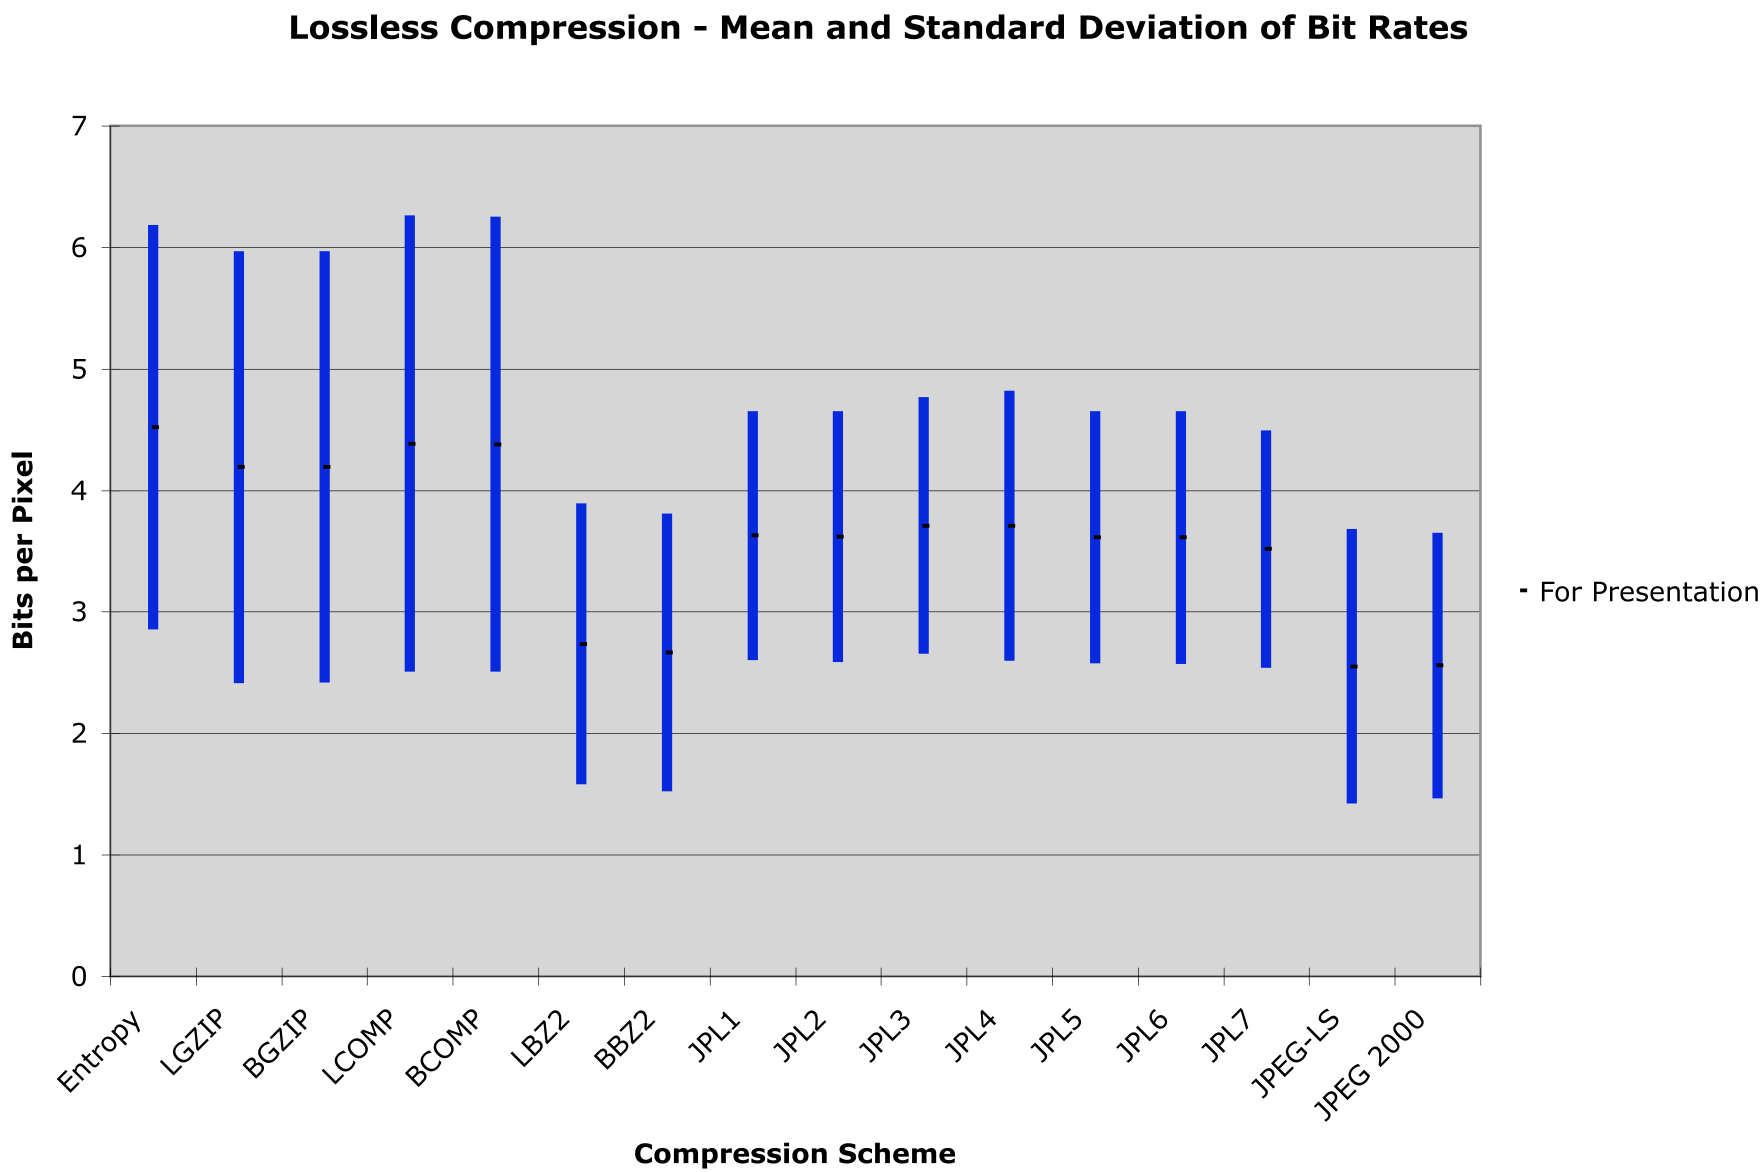

Lossless Compression - Mean and Standard Deviation of Bit Rates
Bits per Pixel
0
1
2
3
4
5
6
7
Entropy
LGZIP
BGZIP
LCOMP
BCOMP
LBZ2
BBZ2
JPL1
JPL2
JPL3
JPL4
JPL5
JPL6
JPL7
JPEG-LS
JPEG 2000
Compression Scheme
For Presentation

Lossless compression

- For Presentation compress better than For Processing - less information
- All compress extremely well - mostly air
- Considerable variation - size of breast ?
- JPEG-LS and JPEG 2000 best
  - Mean CR 6.27 and 6.25 For Presentation
- Lossless JPEG (SV1) poor
  - Mean CR 4.41 For Presentation
  - No run length compression - poor for large areas of air
- Bzip2 does surprisingly well
  - Mean CR 6.00 For Presentation
  - Large block based scheme - knows nothing about images

Variation in compressibility

JPEG-LS Lossless

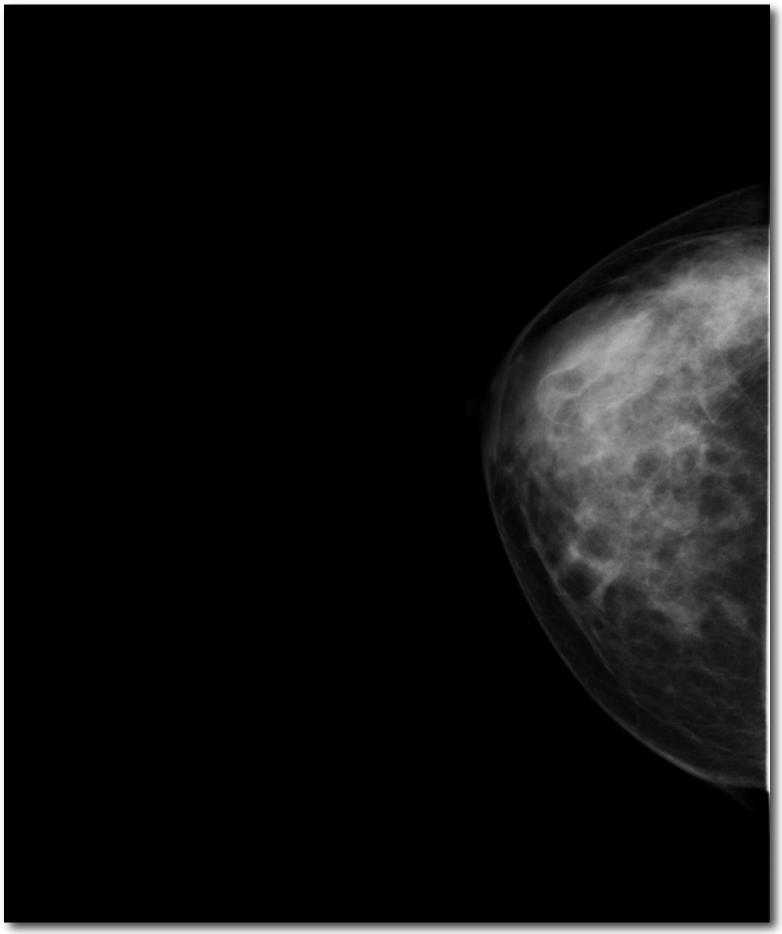

Best - CR 12.9

Worst - CR 3.19

Lossy mammo compression

- Is it OK for any purpose ?
- Are wavelets better than JPEG ?
  - Several experiments suggest not, at compression ratios that are practical
- What compression ratio (bit rate) is OK ?
  - Depends on how much information is in image
  - How much air versus breast
- Region of interest compression
  - Compress background more than breast
  - A feature of many schemes, including J2K

Manufacturer = Fischer
12 in 16 bits per pixel
Matrix = 5625 x 4095
Size = 46,068,750 bytes
Entropy = 4.46 bpp

Lossless J2K = 2.71 bpp
Size = 7,793,250 bytes
CR = 5.91:1

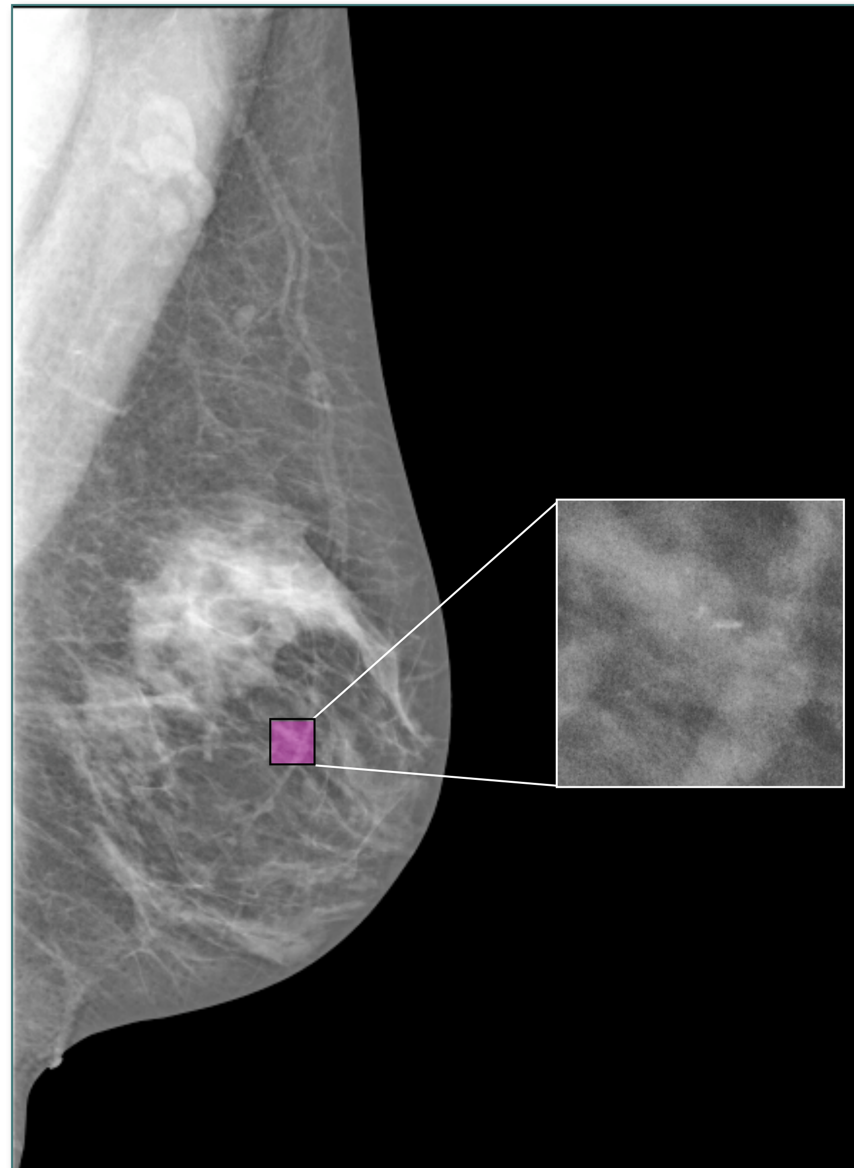

Original
CR 1:1
47MB

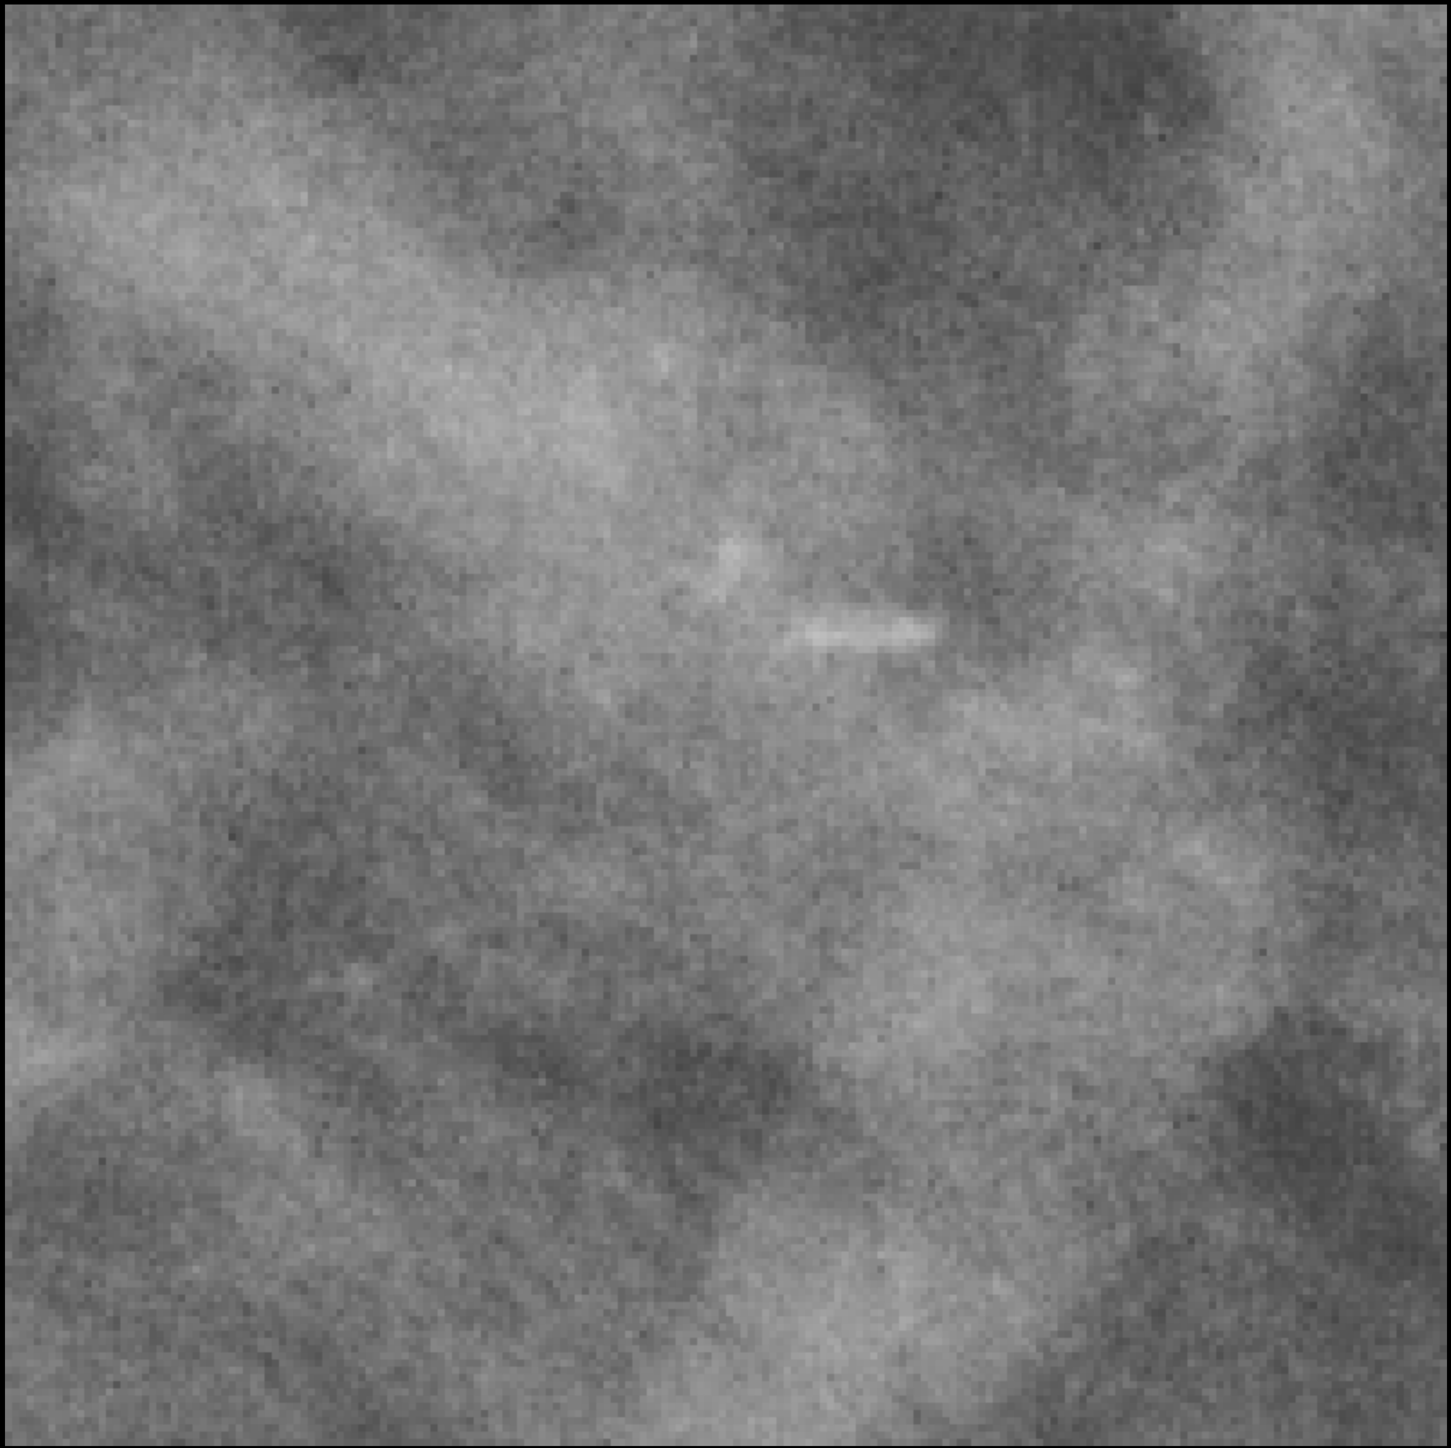

2.0 bpp
CR 8:1
5.7MB
J2K

1.0 bpp
CR 16:1
2.9MB
J2K

0.5 bpp
CR 32:1
1.4MB
J2K

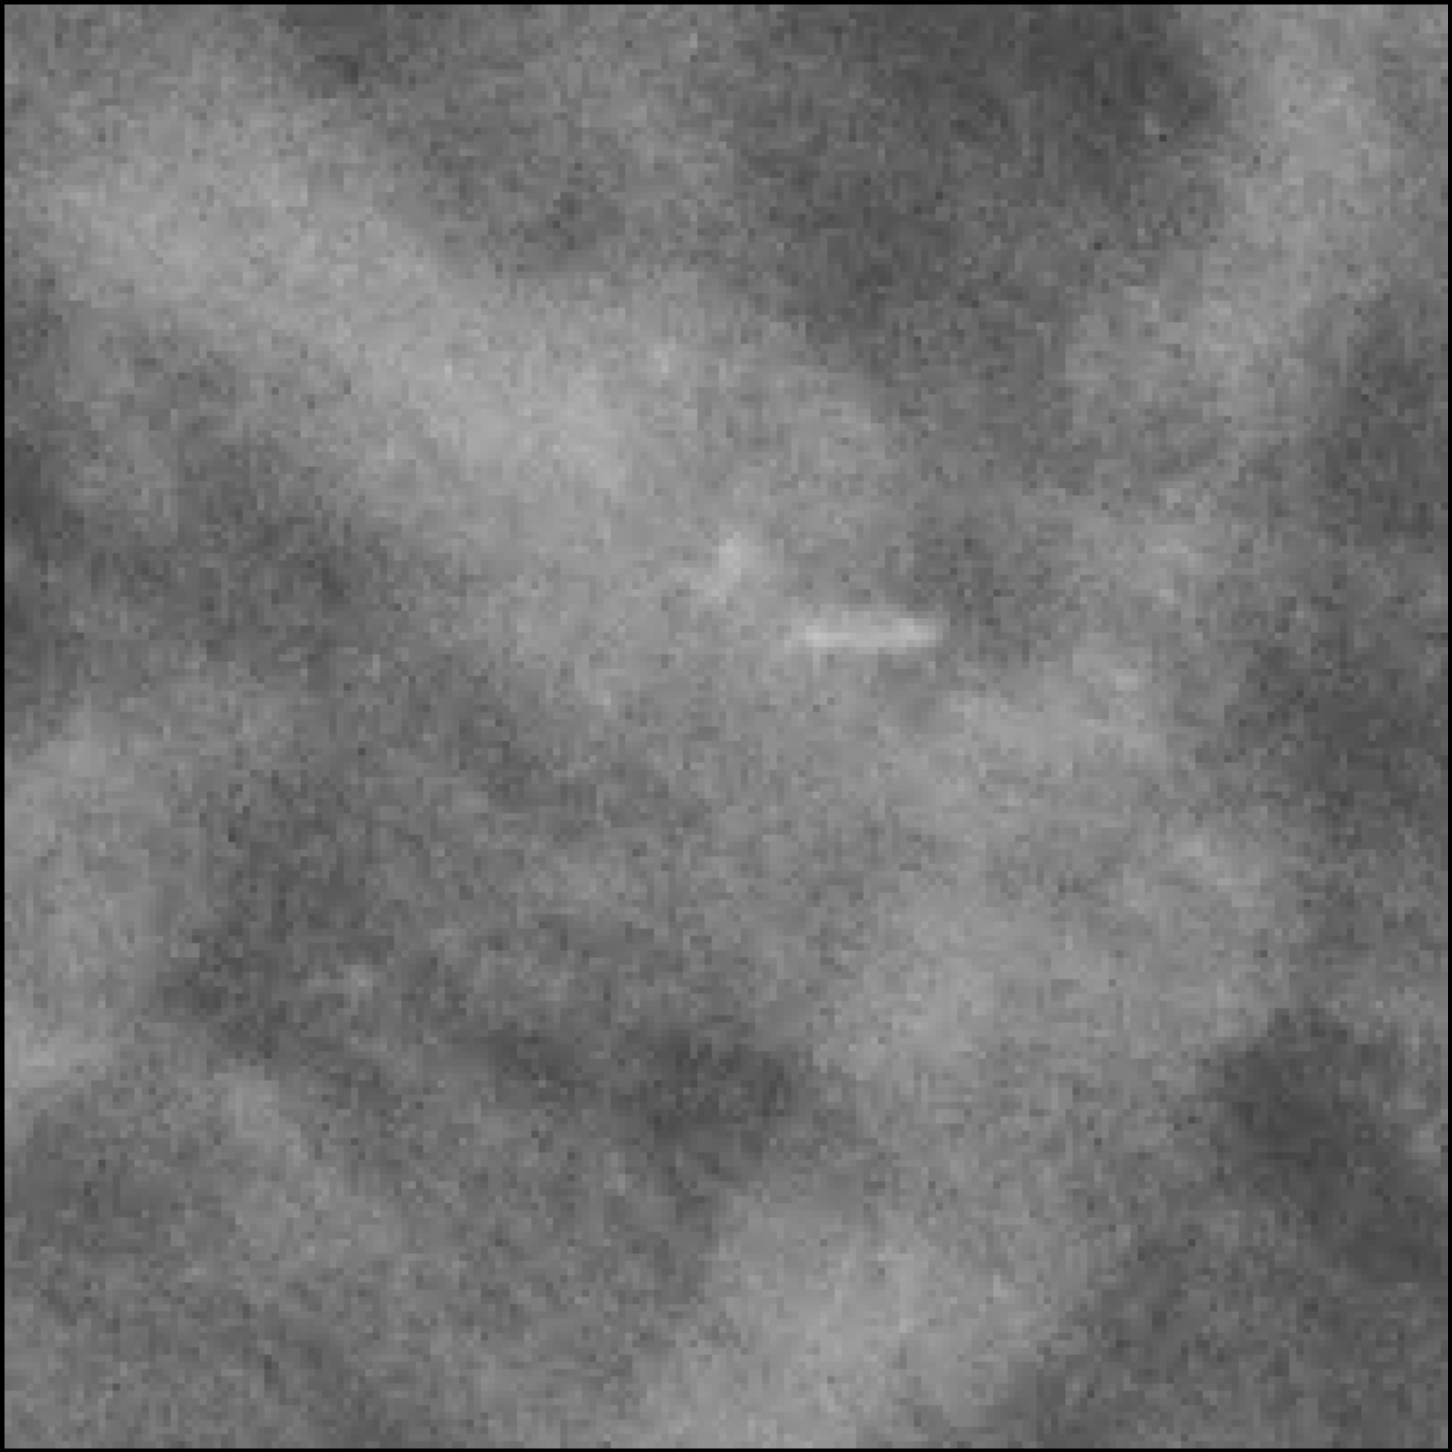

0.375 bpp
CR 43:1
1MB
J2K

0.25 bpp
CR 65:1
710kB
J2K

0.125 bpp
CR 128:1
710kB
J2K

0.375 bpp
CR 43:1
1MB
J2K

0.375 bpp
CR 43:1
1MB
JPEG DCT

Mammography - Studies - I

- Kallergi et al, *Radiology* 2006
  - 500 images, 278 images, 85 cancers
  - digitized film (not FFDM)
    - 60 µm 14 bit specialized digitizer
    - cropped to breast size
  - proprietary adaptive wavelet scheme (not JPEG 2000)
  - observer-performance
    - 3 readers
    - localization ROC
      - likelihood of malignancy 5-point scale
      - 200 pixel radius
    - 5MP 8-bit CRT, calibrated (DICOM GSDF), with zoom

Mammography - Studies - II

- Kallergi et al, *Radiology* 2006 (cont'd)
  - three combinations
    - normal vs. malignant
    - benign vs. malignant
    - normal + benign vs. malignant
  - compression rates varied per image
    - scheme compresses until quality metric satisfied
    - 14:1 to 2051:1; mean about 55:1; 60% > 100:1
  - significant differences ($p < 0.05$) in AROC & ALROC
    - all readers, most combinations
    - did ***<u>BETTER</u>*** on compressed images !
    - postulated to be due to de-noising (smoothing)

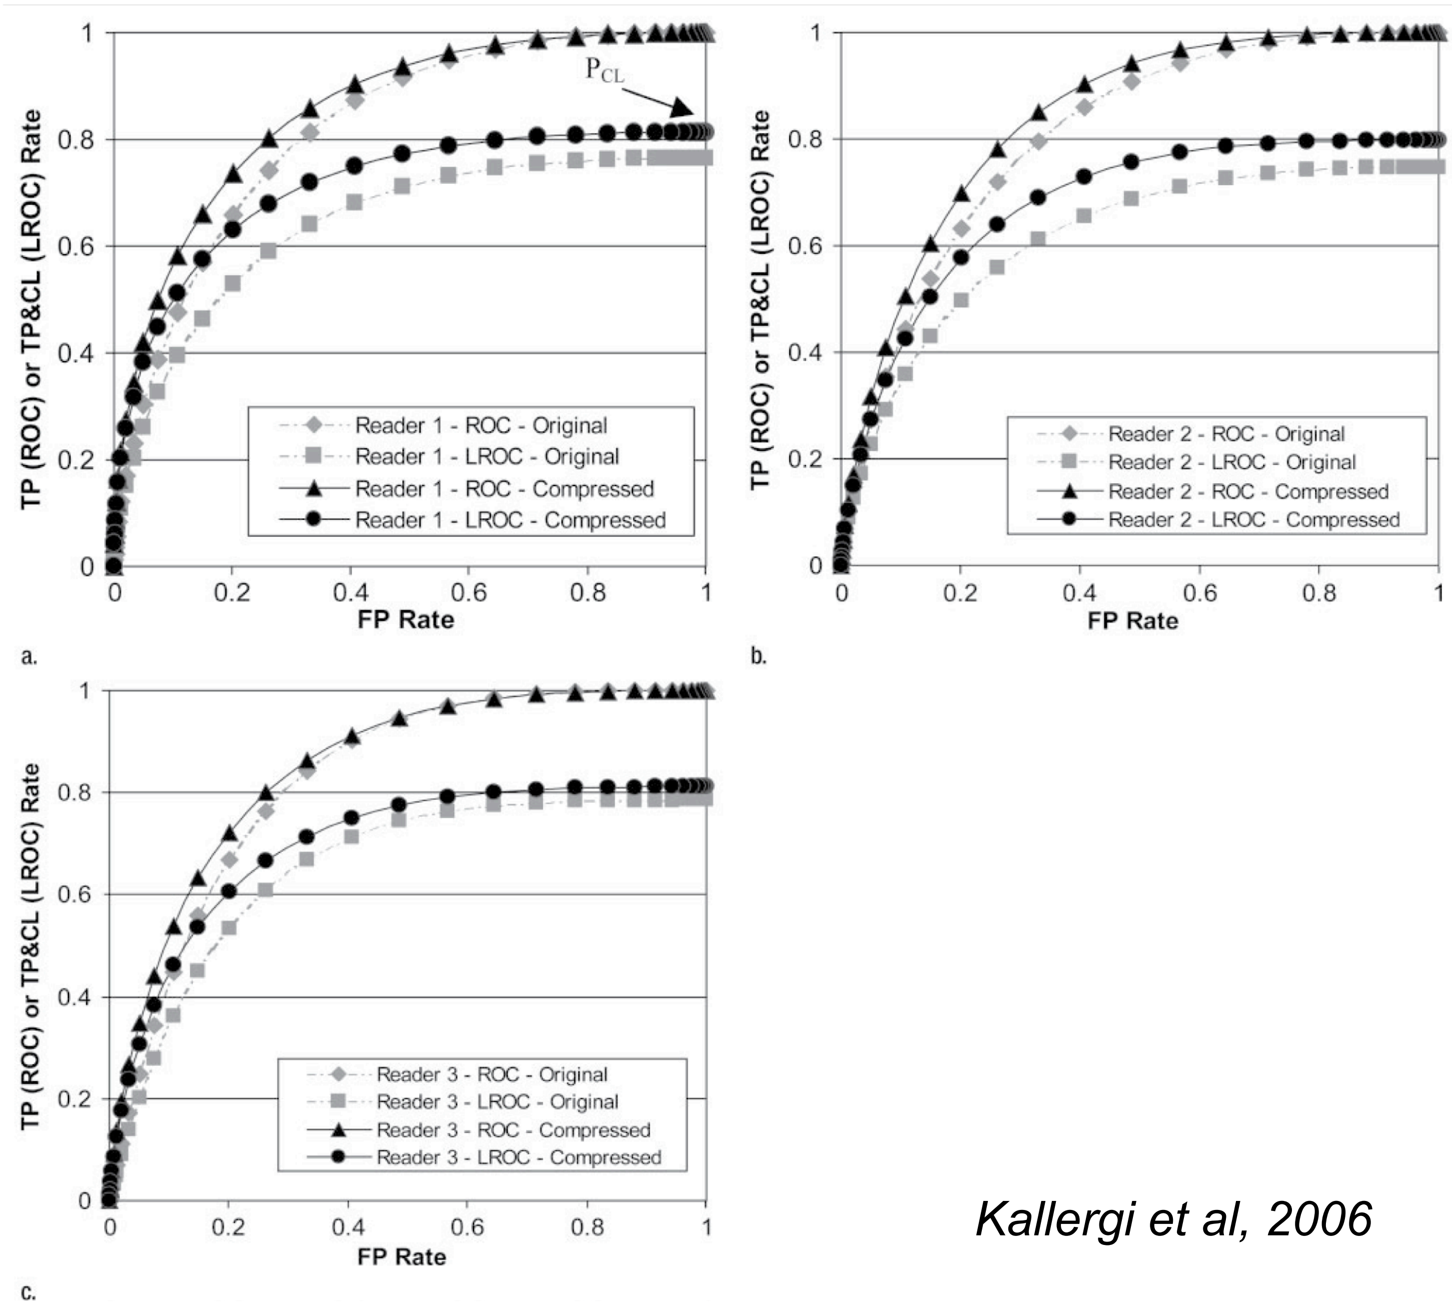

Kallergi et al, 2006

Mammography - Studies - III

- Implications of Kallergi article
  - for study of standard compression schemes
    - crop to breast first ?
    - fixed vs. adaptive compression rate for JPEG 2000
  - digital versus digitized mammograms
    - different noise characteristics
    - different bit depth
    - different pixel size
  - choice of methodology
    - ROC, LROC, FROC … others like agreement
  - power of future studies
    - a difference was found (just unexpected direction)
    - statistically significant but was it important clinically
    - choice of number of cases/readers was not described

Mammography - Studies - IV

- Penedo et al, *Radiology* 2005
  - 112 images, 60 patients
  - digitized film (not FFDM)
    - 50 µm 12 bit commercial digitizer
    - cropped to breast size
  - proprietary wavelet scheme (SPIHT) and JPEG 2000
  - observer-performance
    - 5 readers
    - free response ROC detection of masses & micro-calcifications
    - printed to film for interpretation
  - no difference between originals & compressed 40:1 & 80:1
  - 95% confidence intervals include 0 and within 80% power

Multi-frame & 3D compression

- Original CT and MR SOP Classes are single frame
  - Compression only possible within a single frame
  - Lossless - typically 3:1 or 4:1 for CT and MR
- Multi-frame objects
  - Opportunity to take advantage of redundancy between frames
  - Spatial redundancy - JPEG 2000 Part 2
    - Lossless gain modest, lossy gain more substantial
  - Motion prediction - MPEG-2 and others
  - New schemes - H.264/MPEG-4 Part 10
  - Entire dataset (e.g., 3D volume) or adjacent slabs

Single frame lossless compression

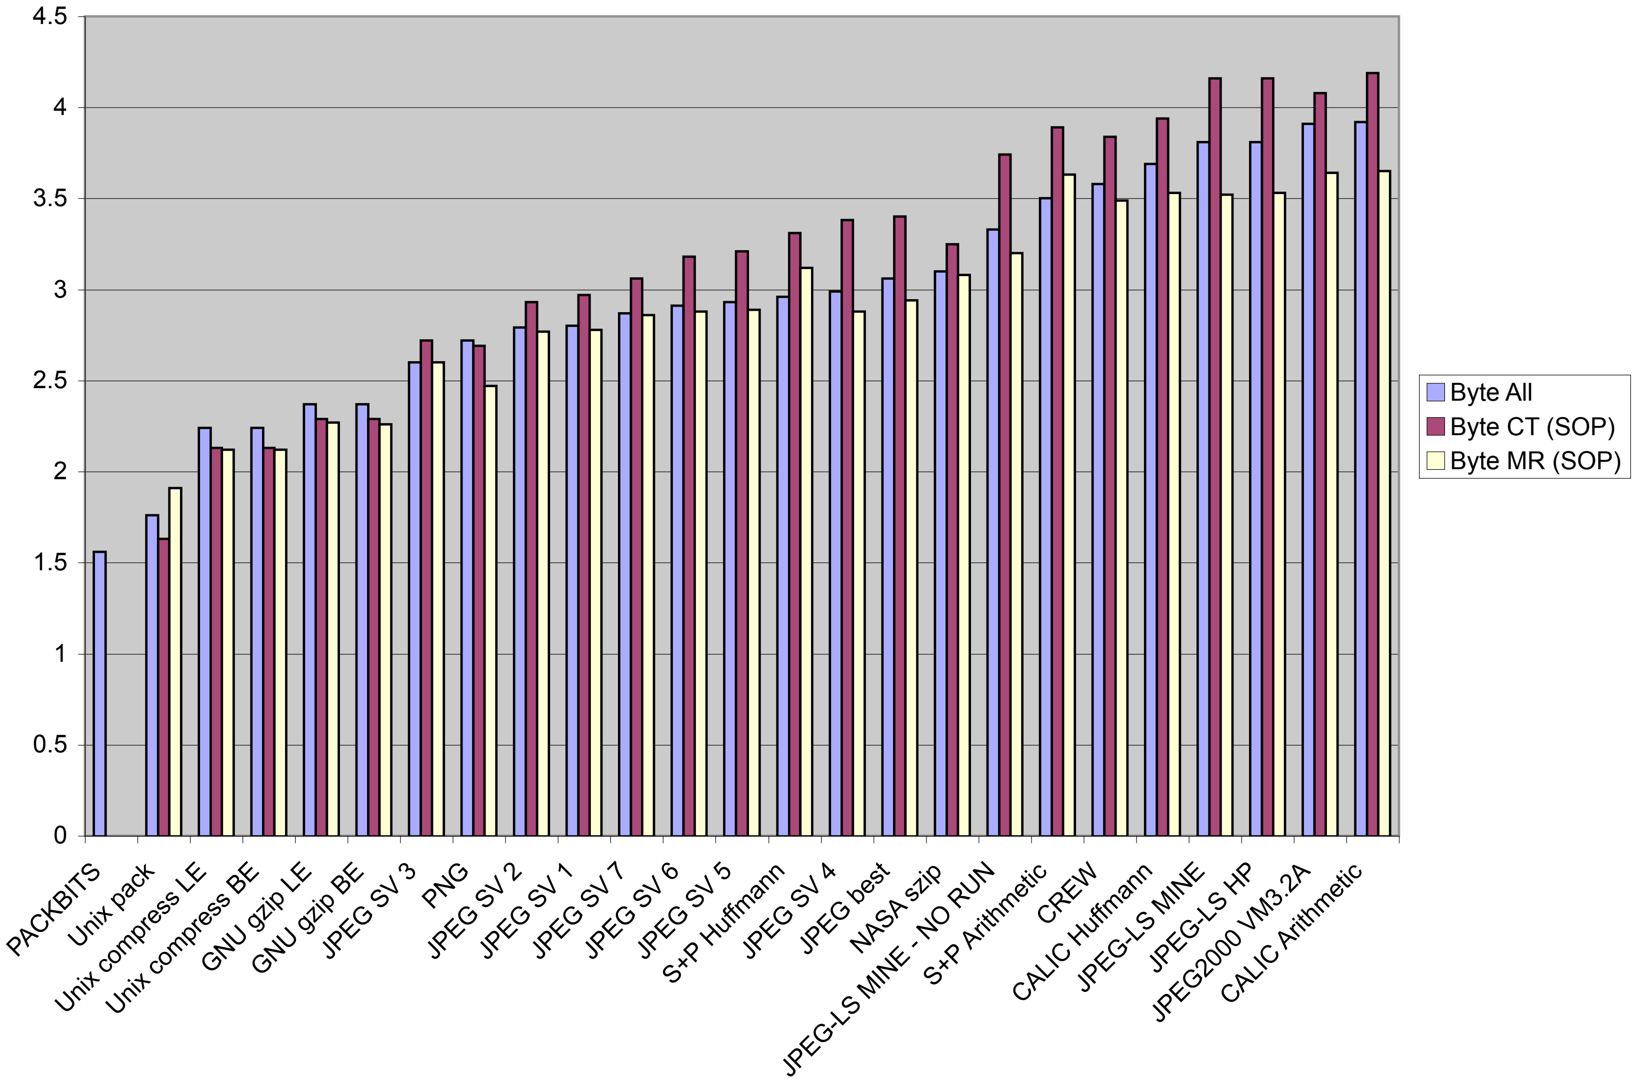

Lossless JPEG 2000 Compression (Alexis Tzannes, Aware, 2003)

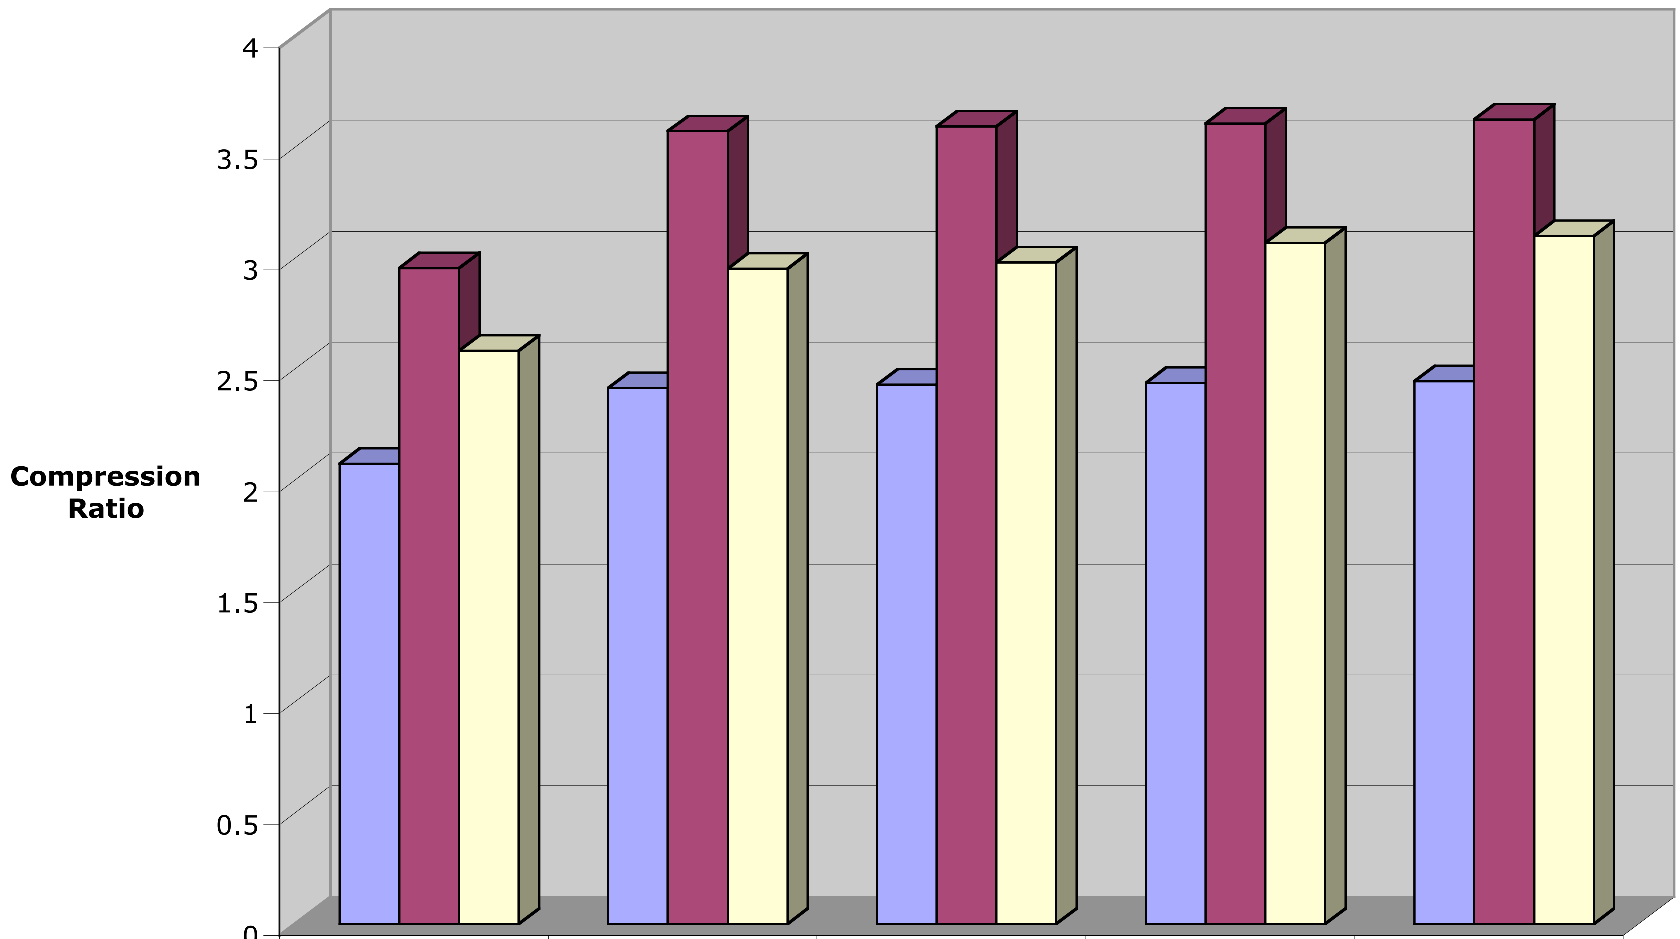

	single	20	40	80	all
127x256x8 7.9MB	2.073490814	2.415902141	2.430769231	2.438271605	2.445820433
449x512x16 224MB	2.955145119	3.572567783	3.595505618	3.607085346	3.624595469
620x512x16 310MB	2.583333333	2.952380952	2.980769231	3.069306931	3.1

Slices in 3rd dimension

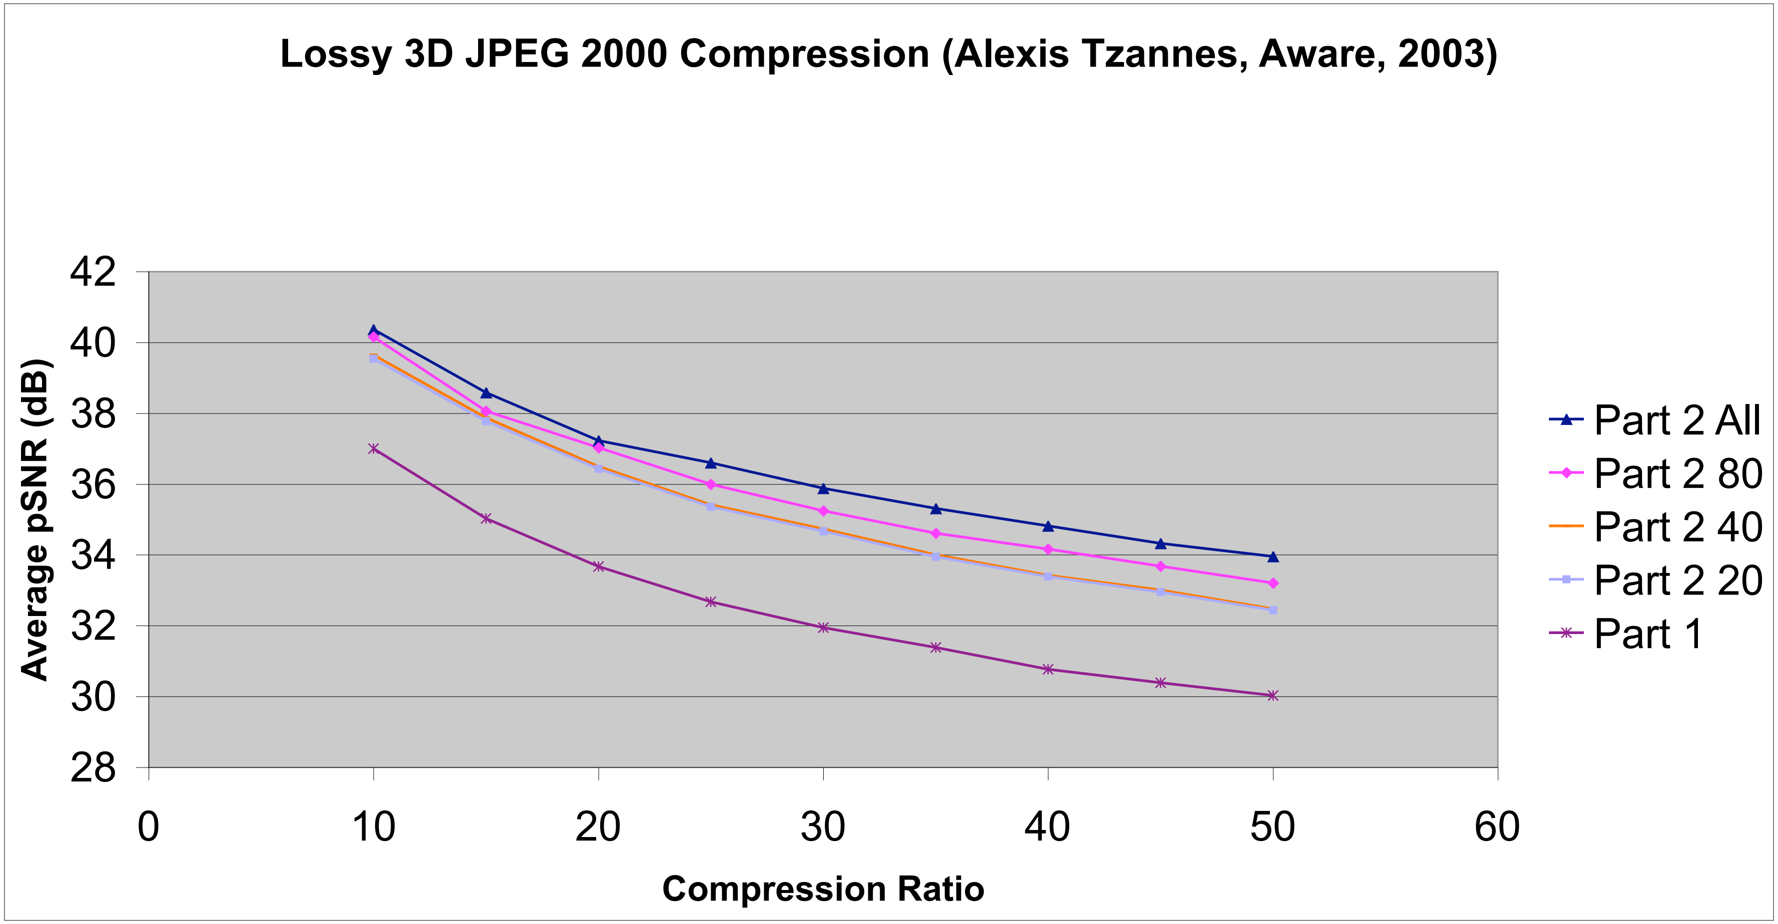

Lossy 3D JPEG 2000 Compression (Alexis Tzannes, Aware, 2003)
Average pSNR (dB)
42
40
38
36
34
32
30
28
0
10
20
30
40
50
60
Compression Ratio
Part 2 All
Part 2 80
Part 2 40
Part 2 20
Part 1

2D JPEG 2000 0.625mm slices

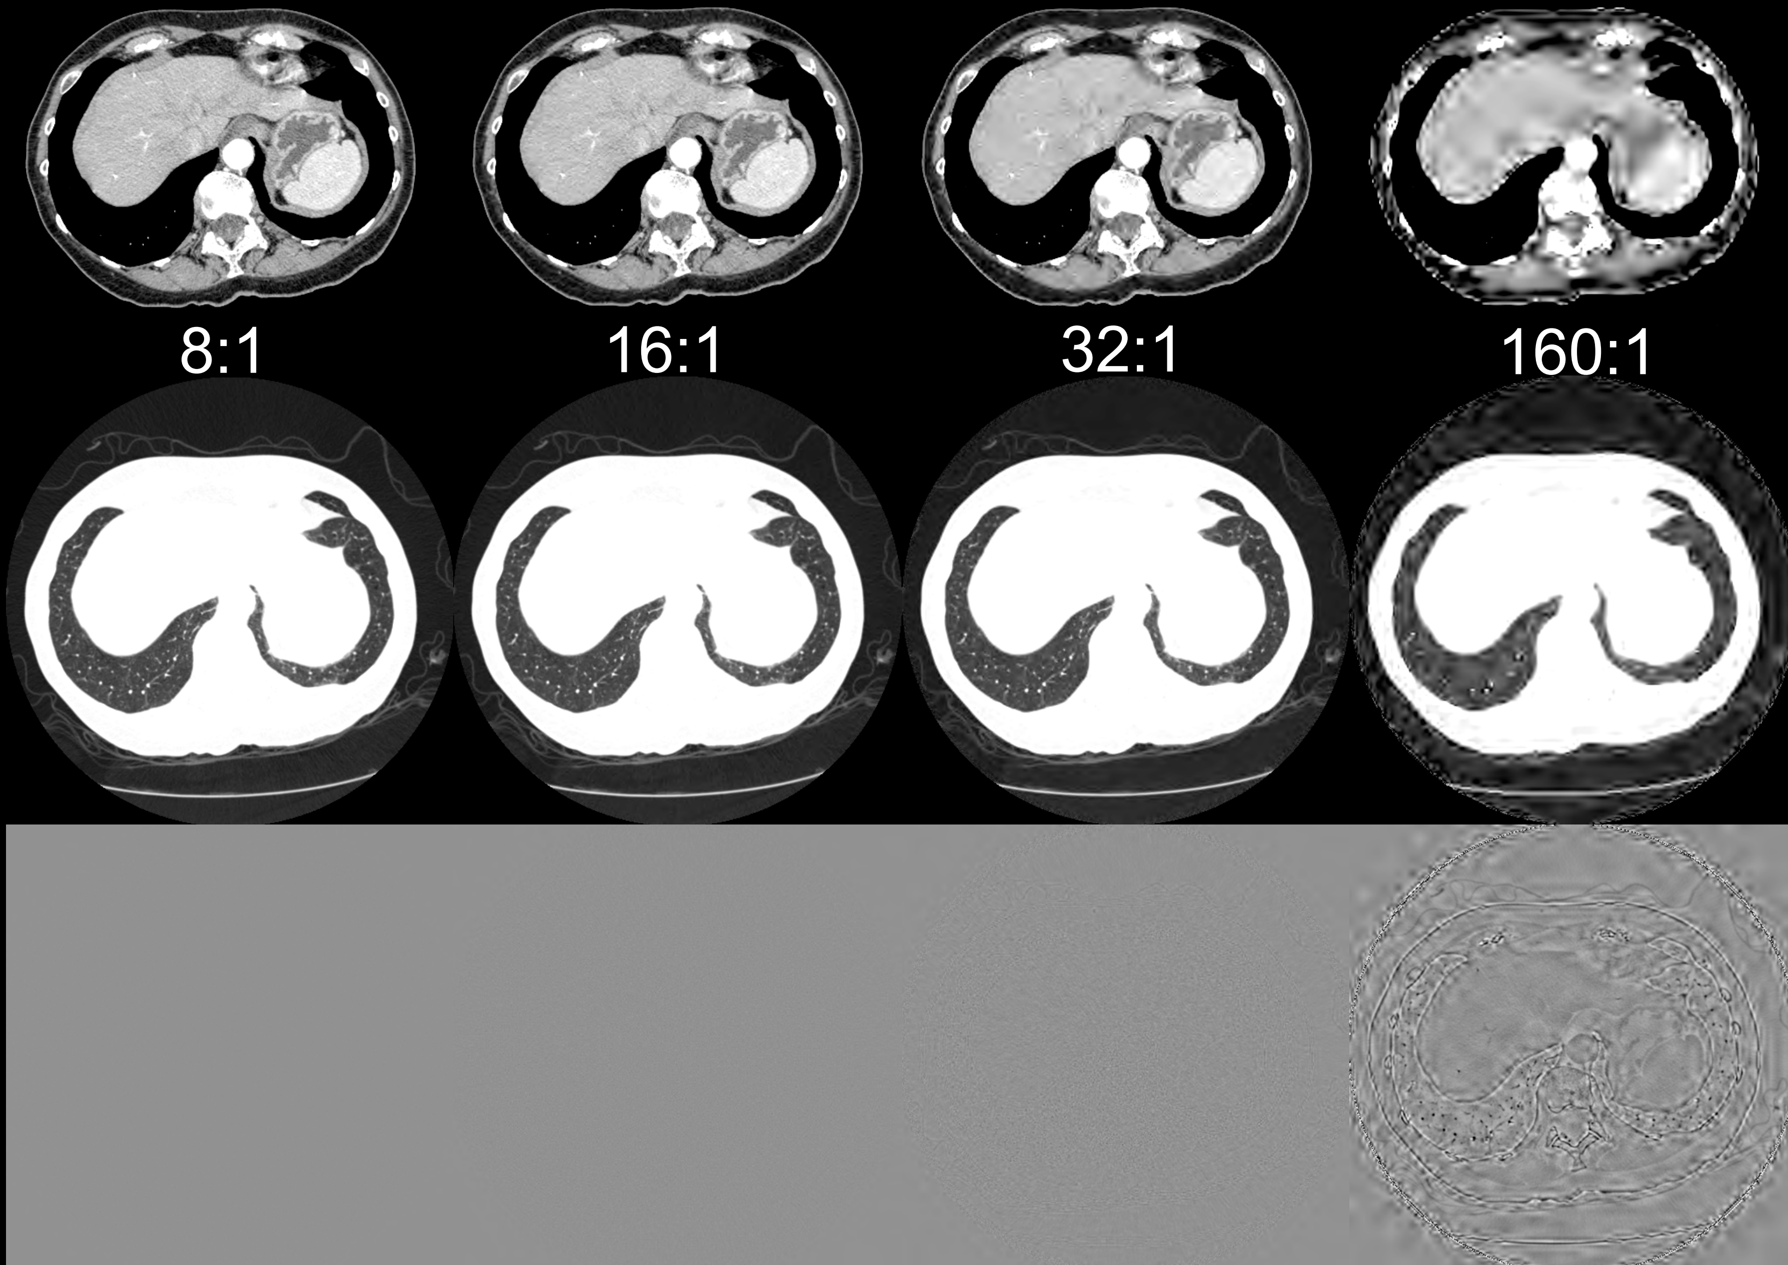

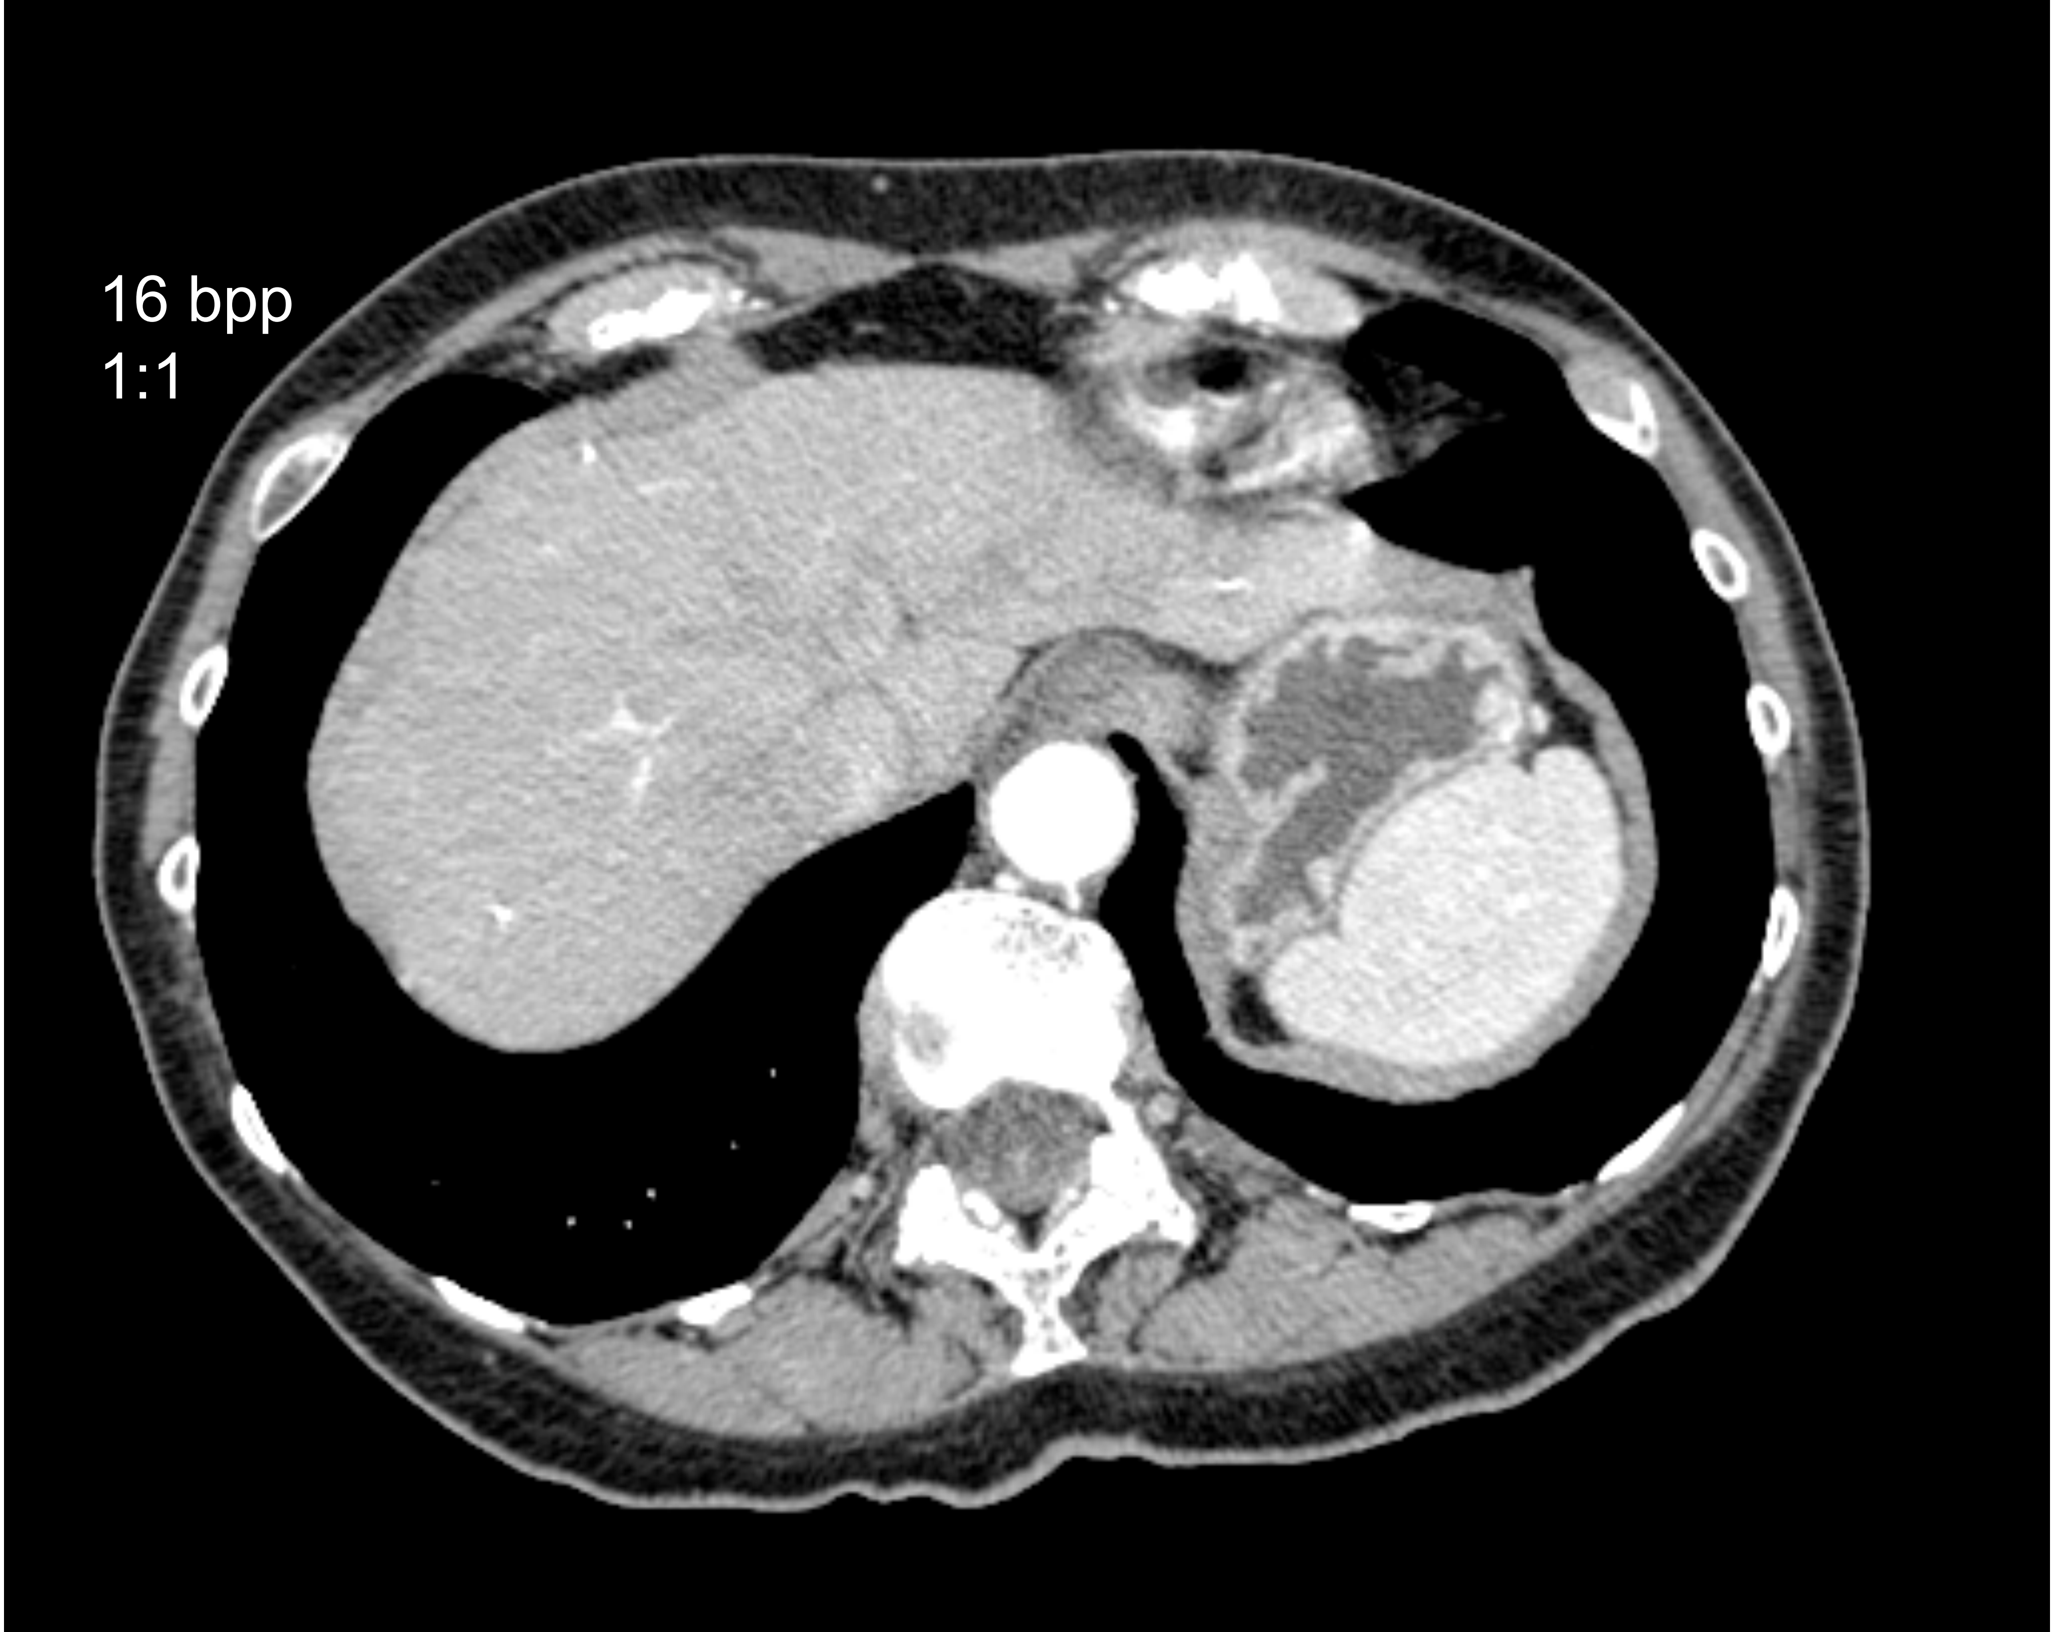
16 bpp
1:1

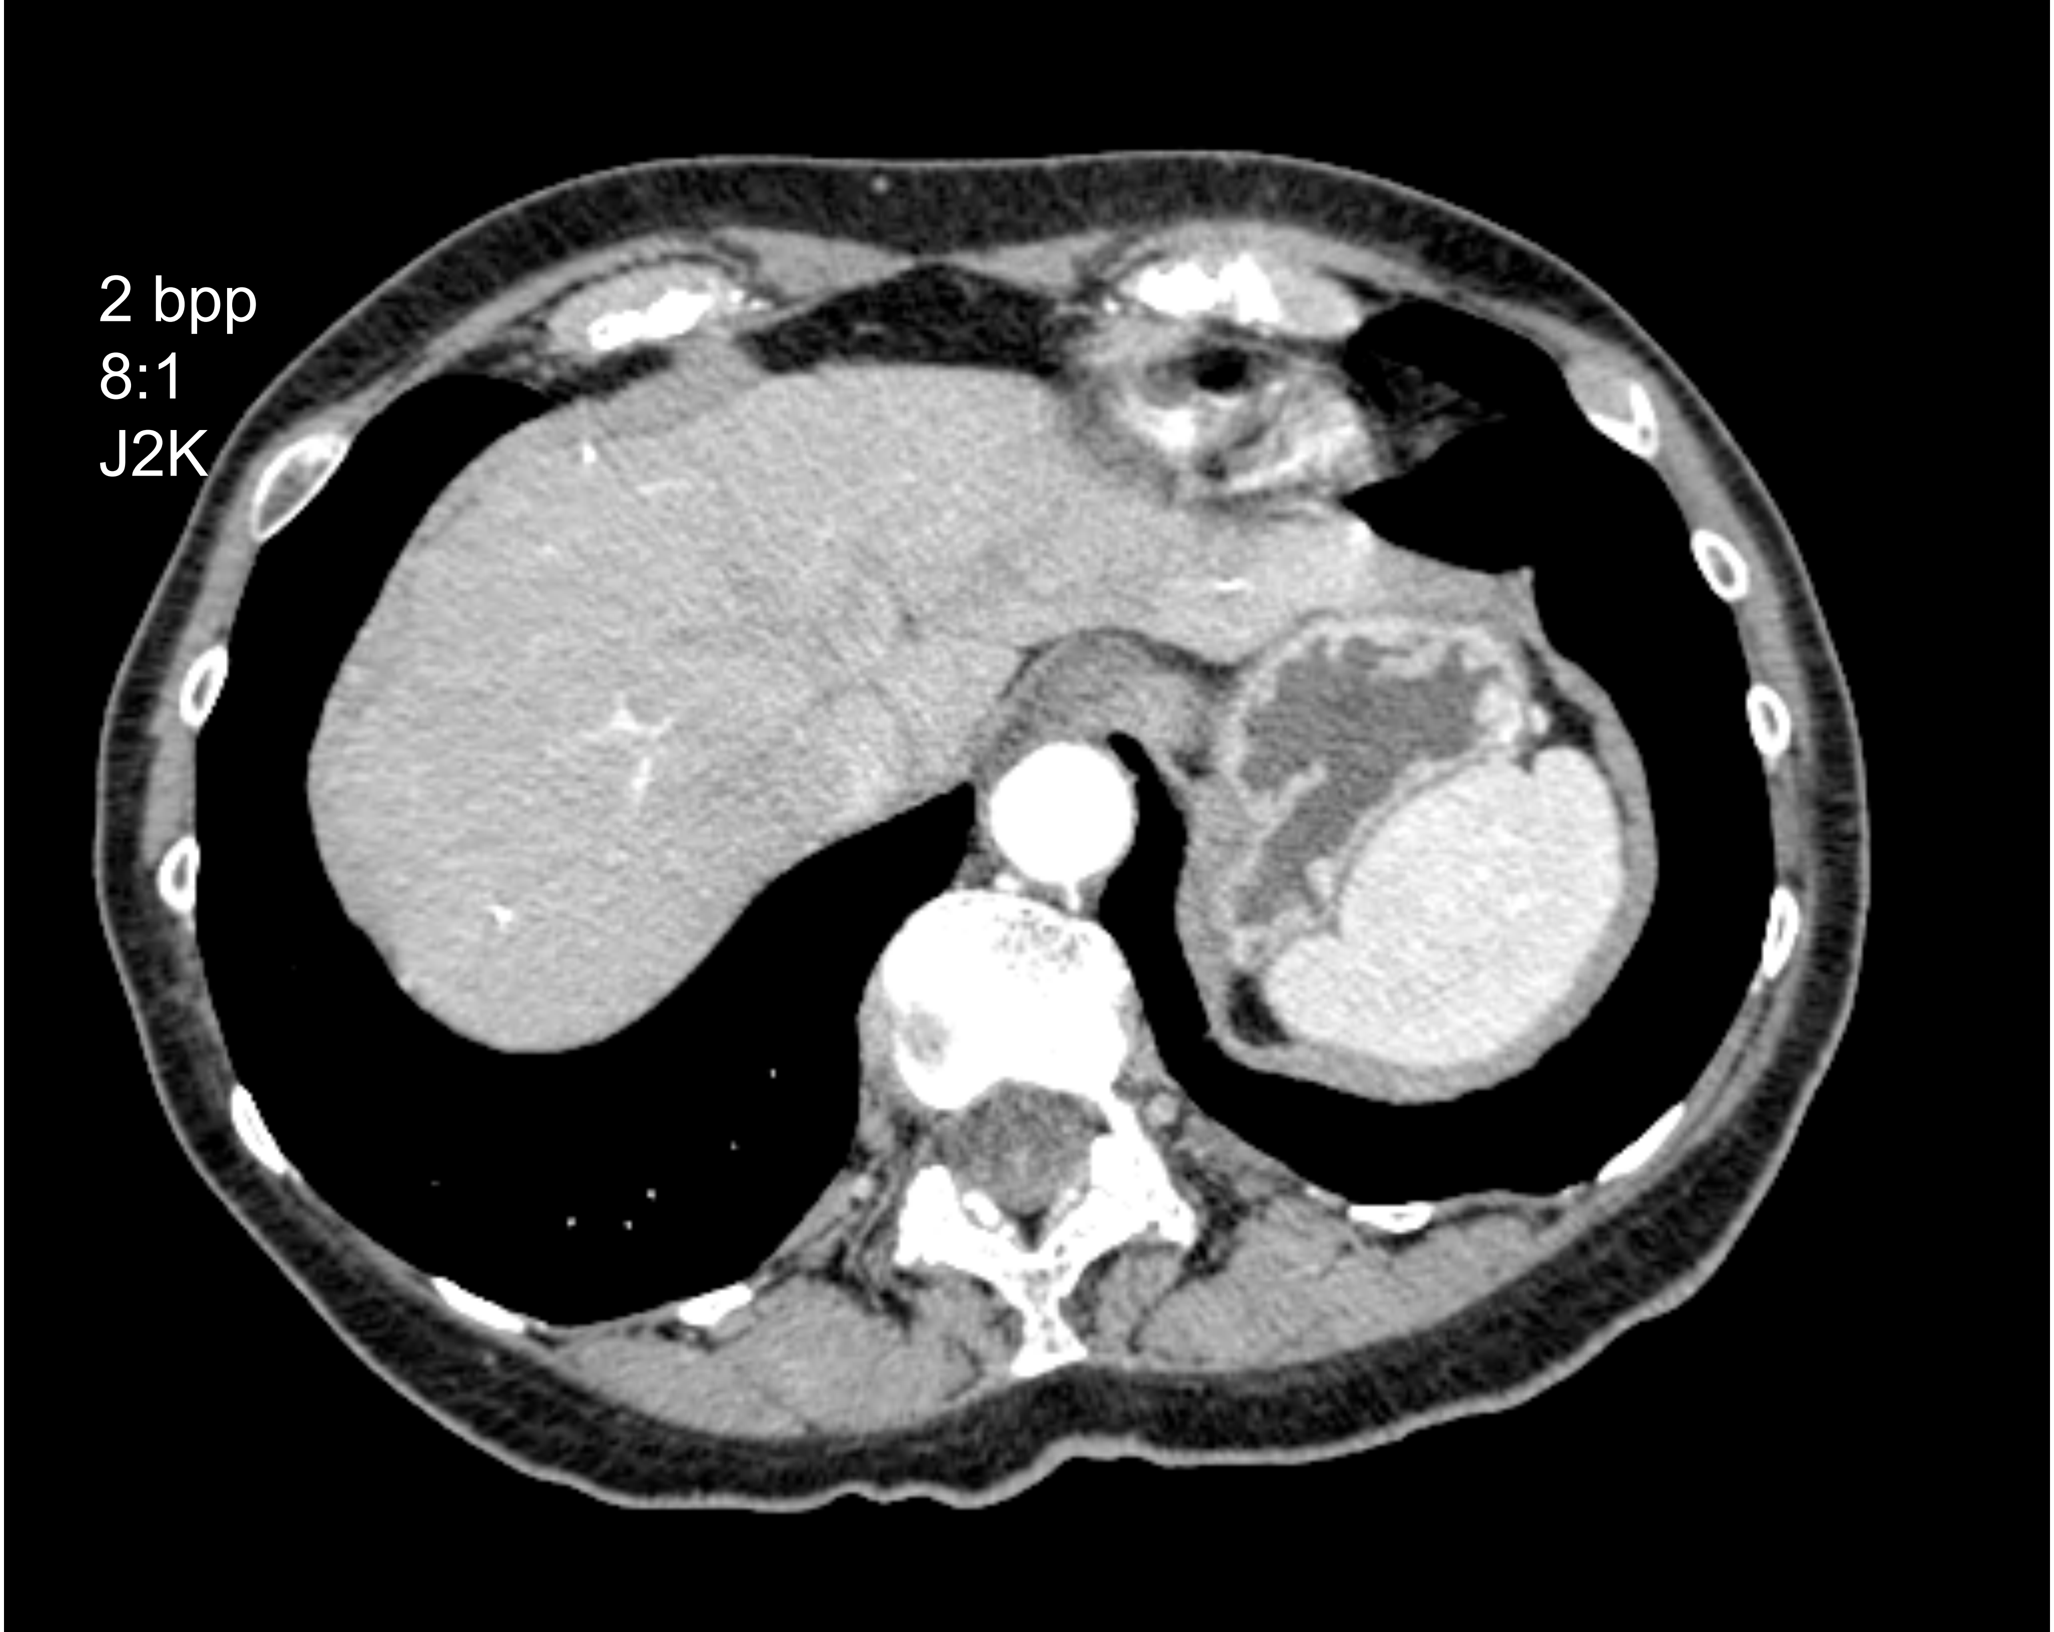
2 bpp
8:1
J2K

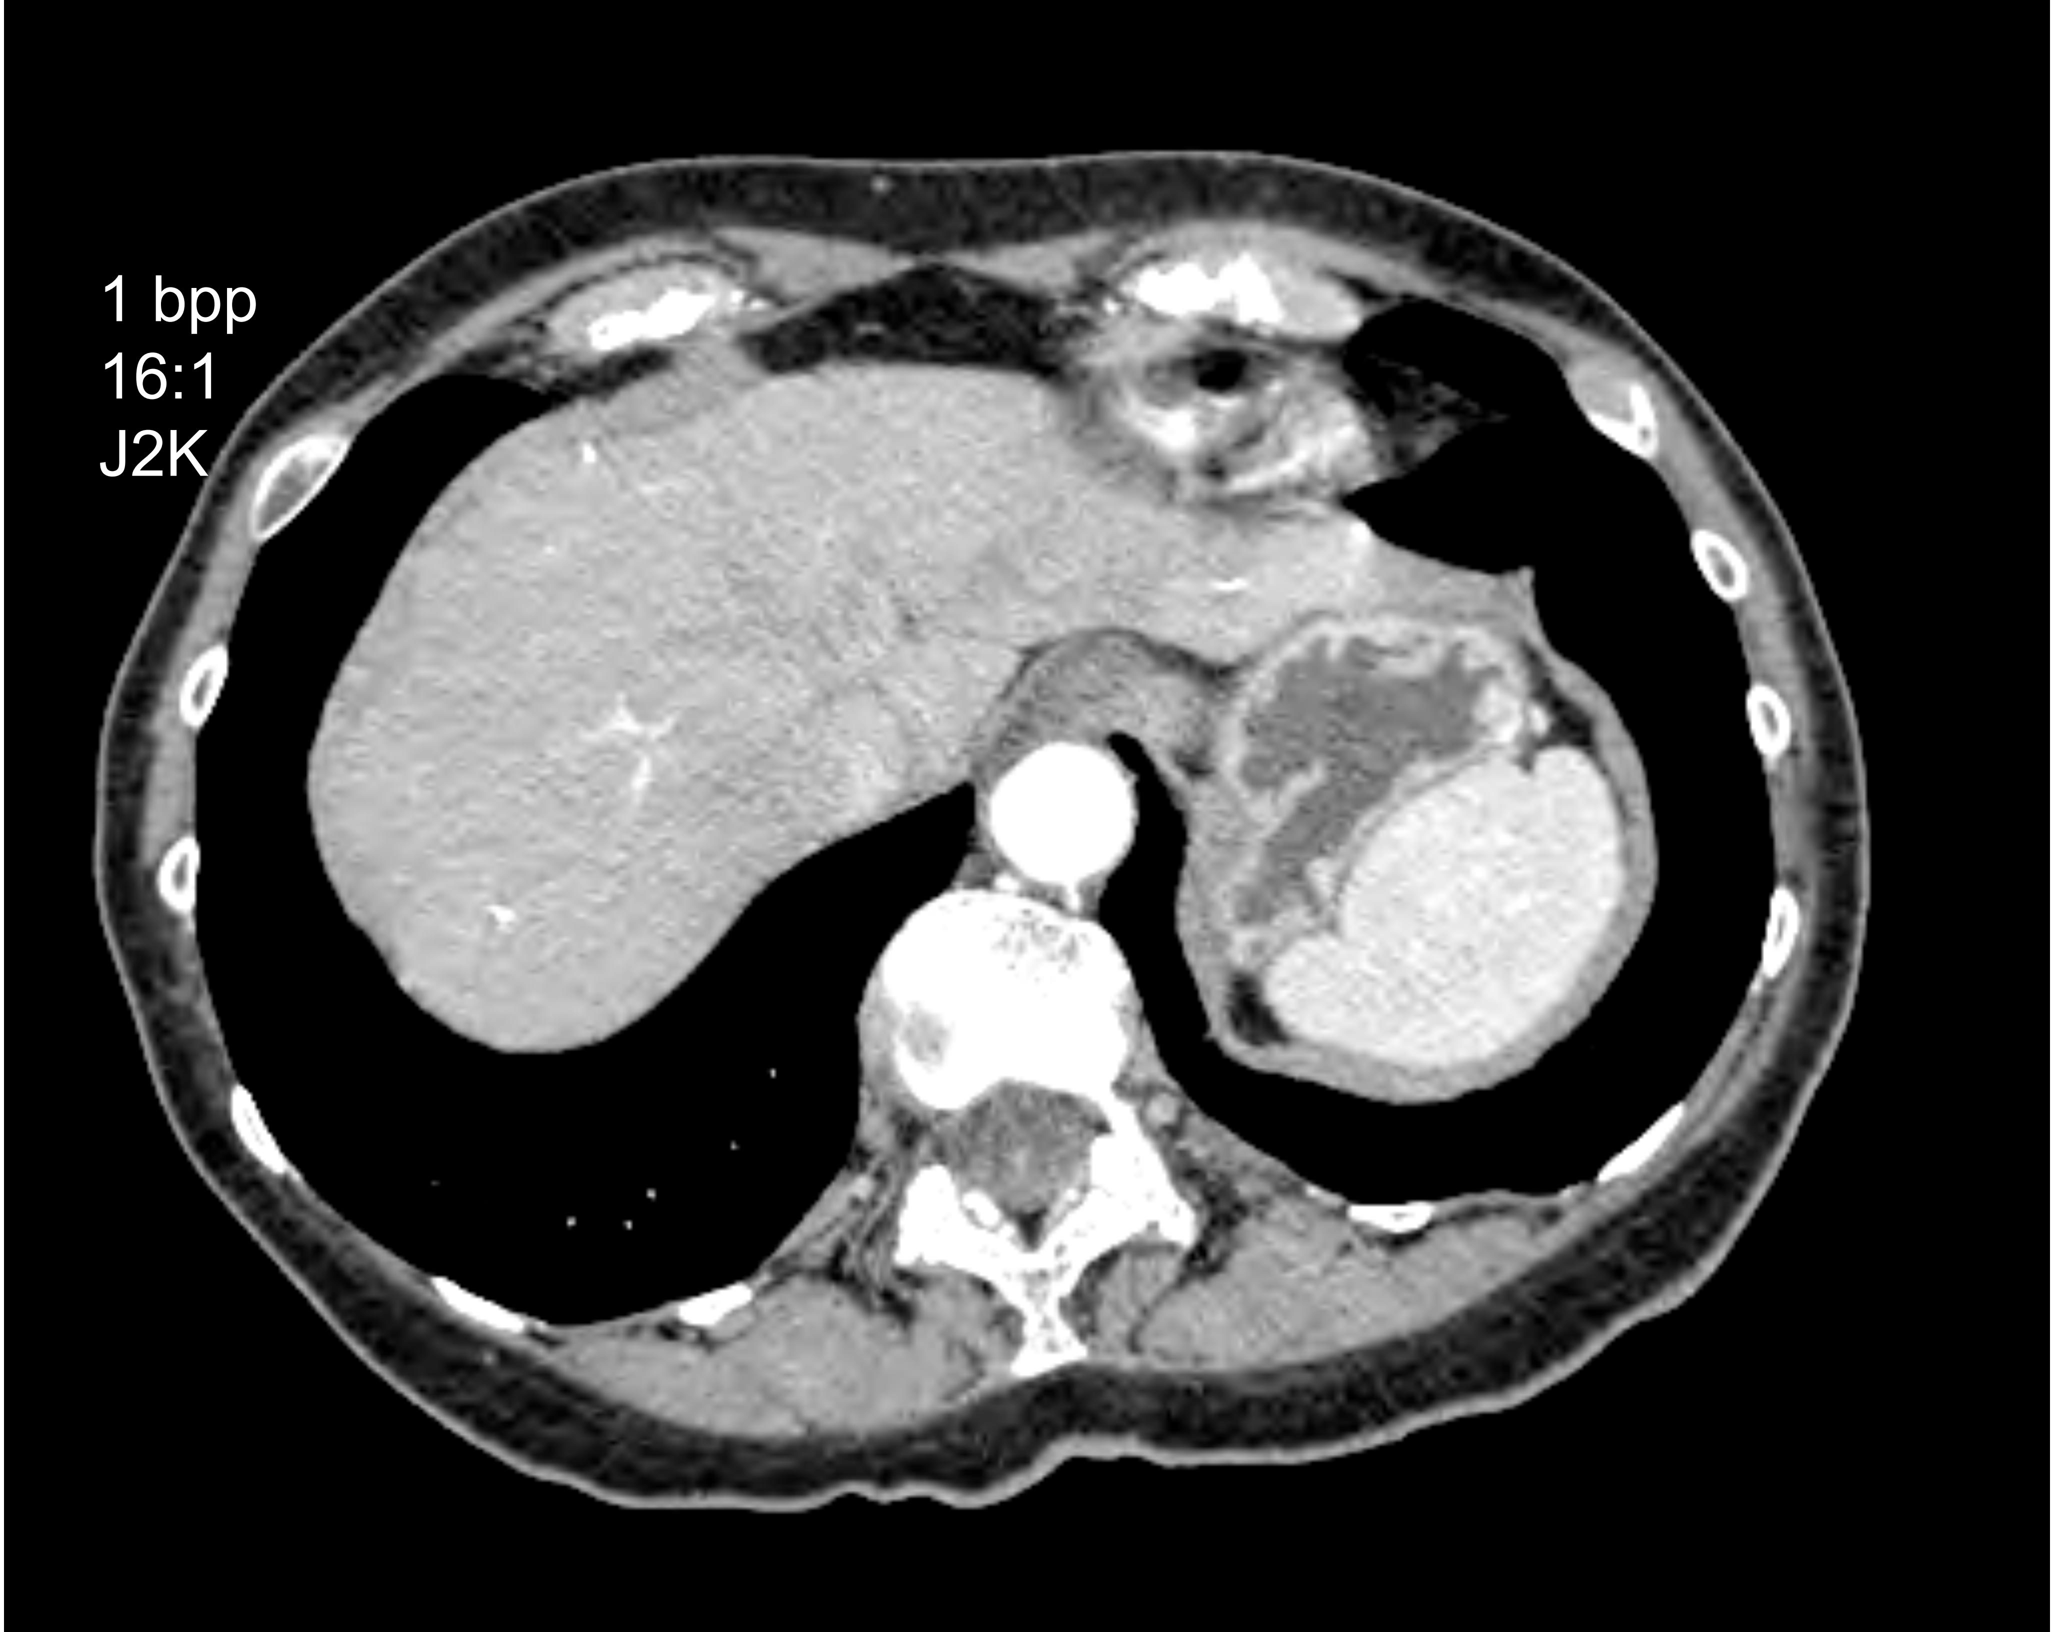
1 bpp
16:1
J2K

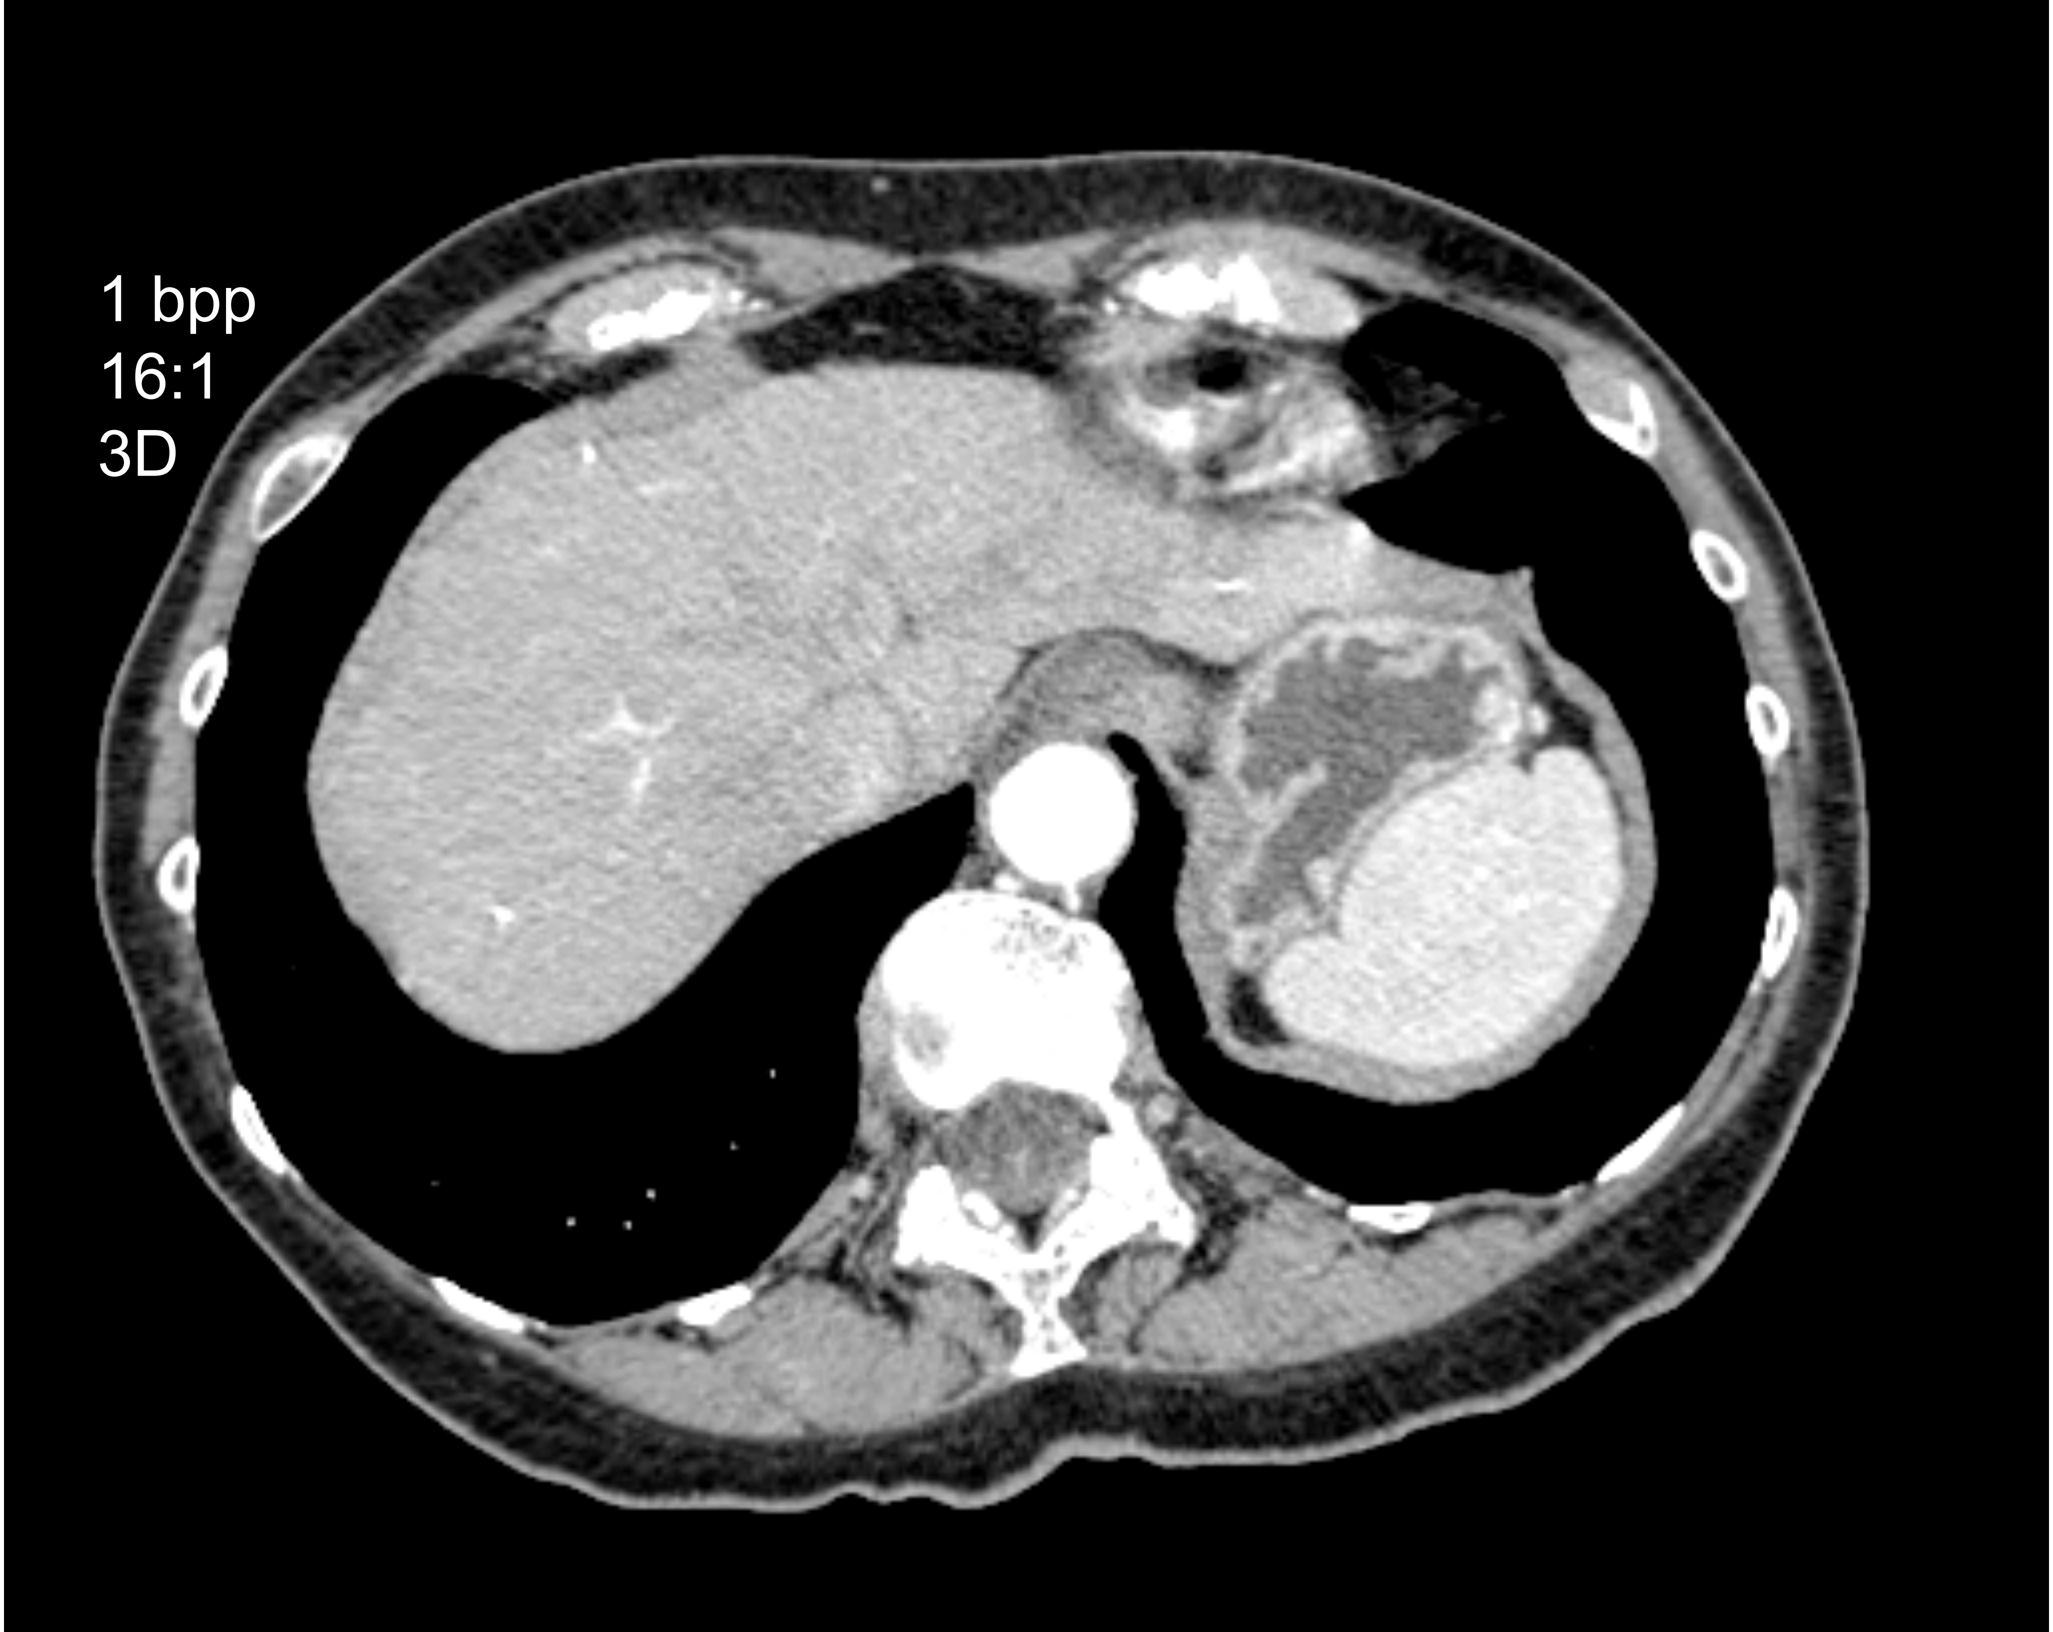
1 bpp
16:1
3D

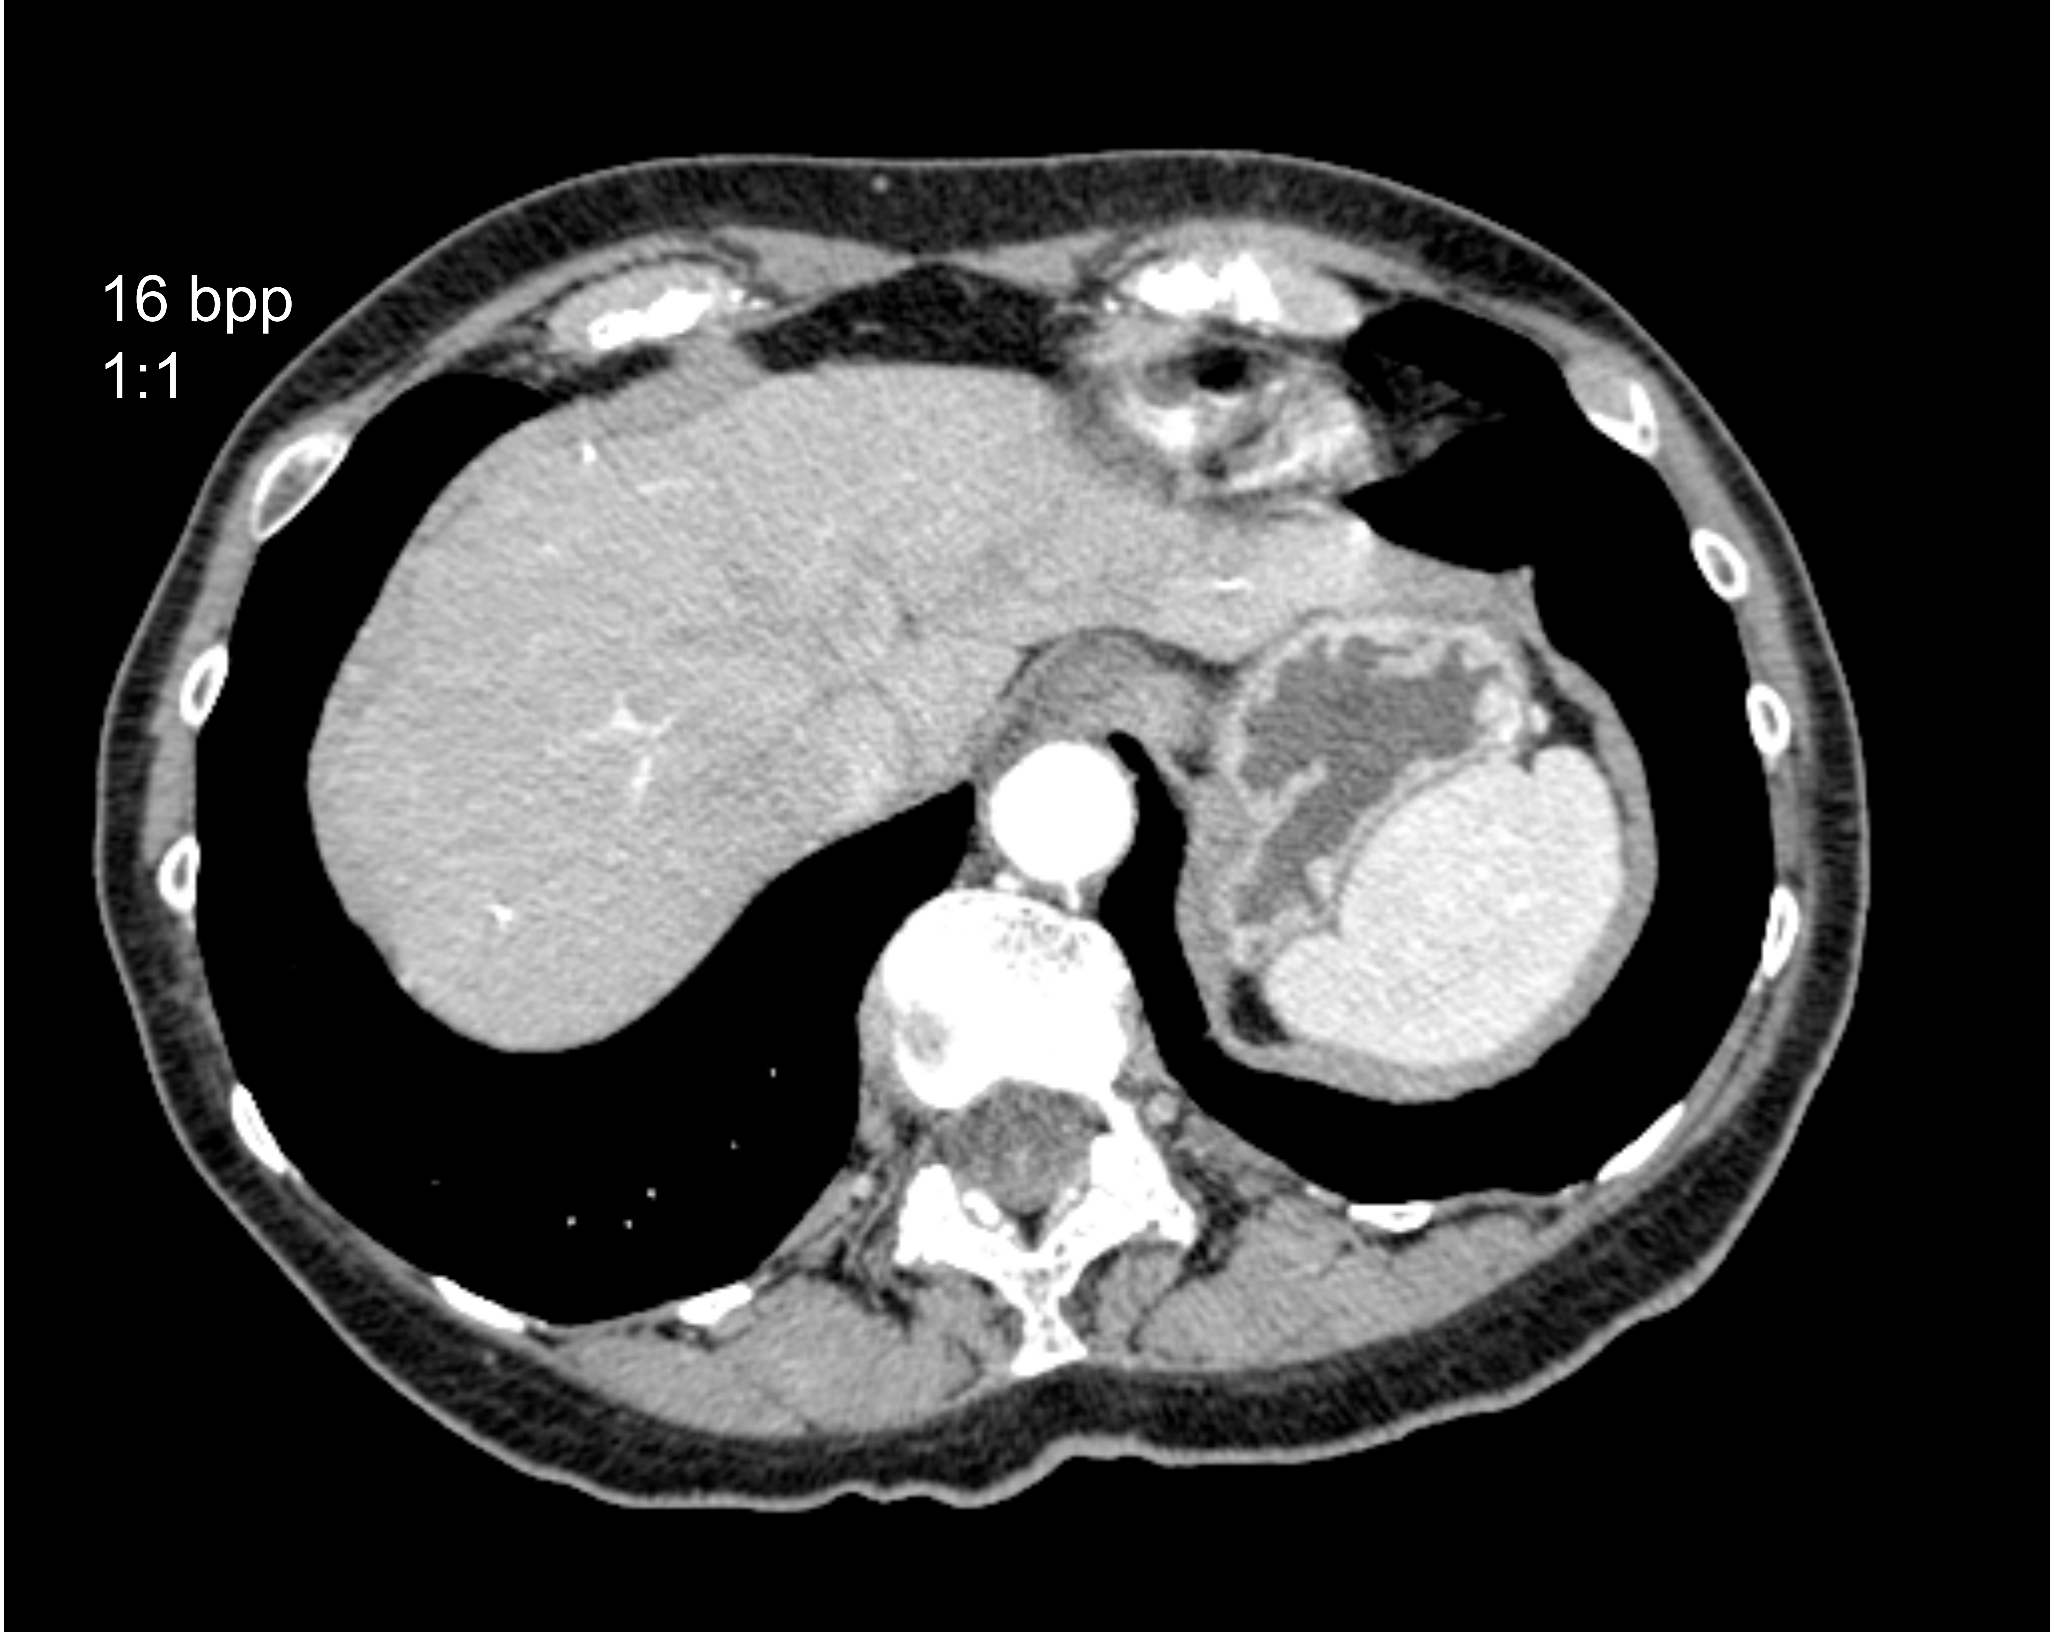
16 bpp
1:1

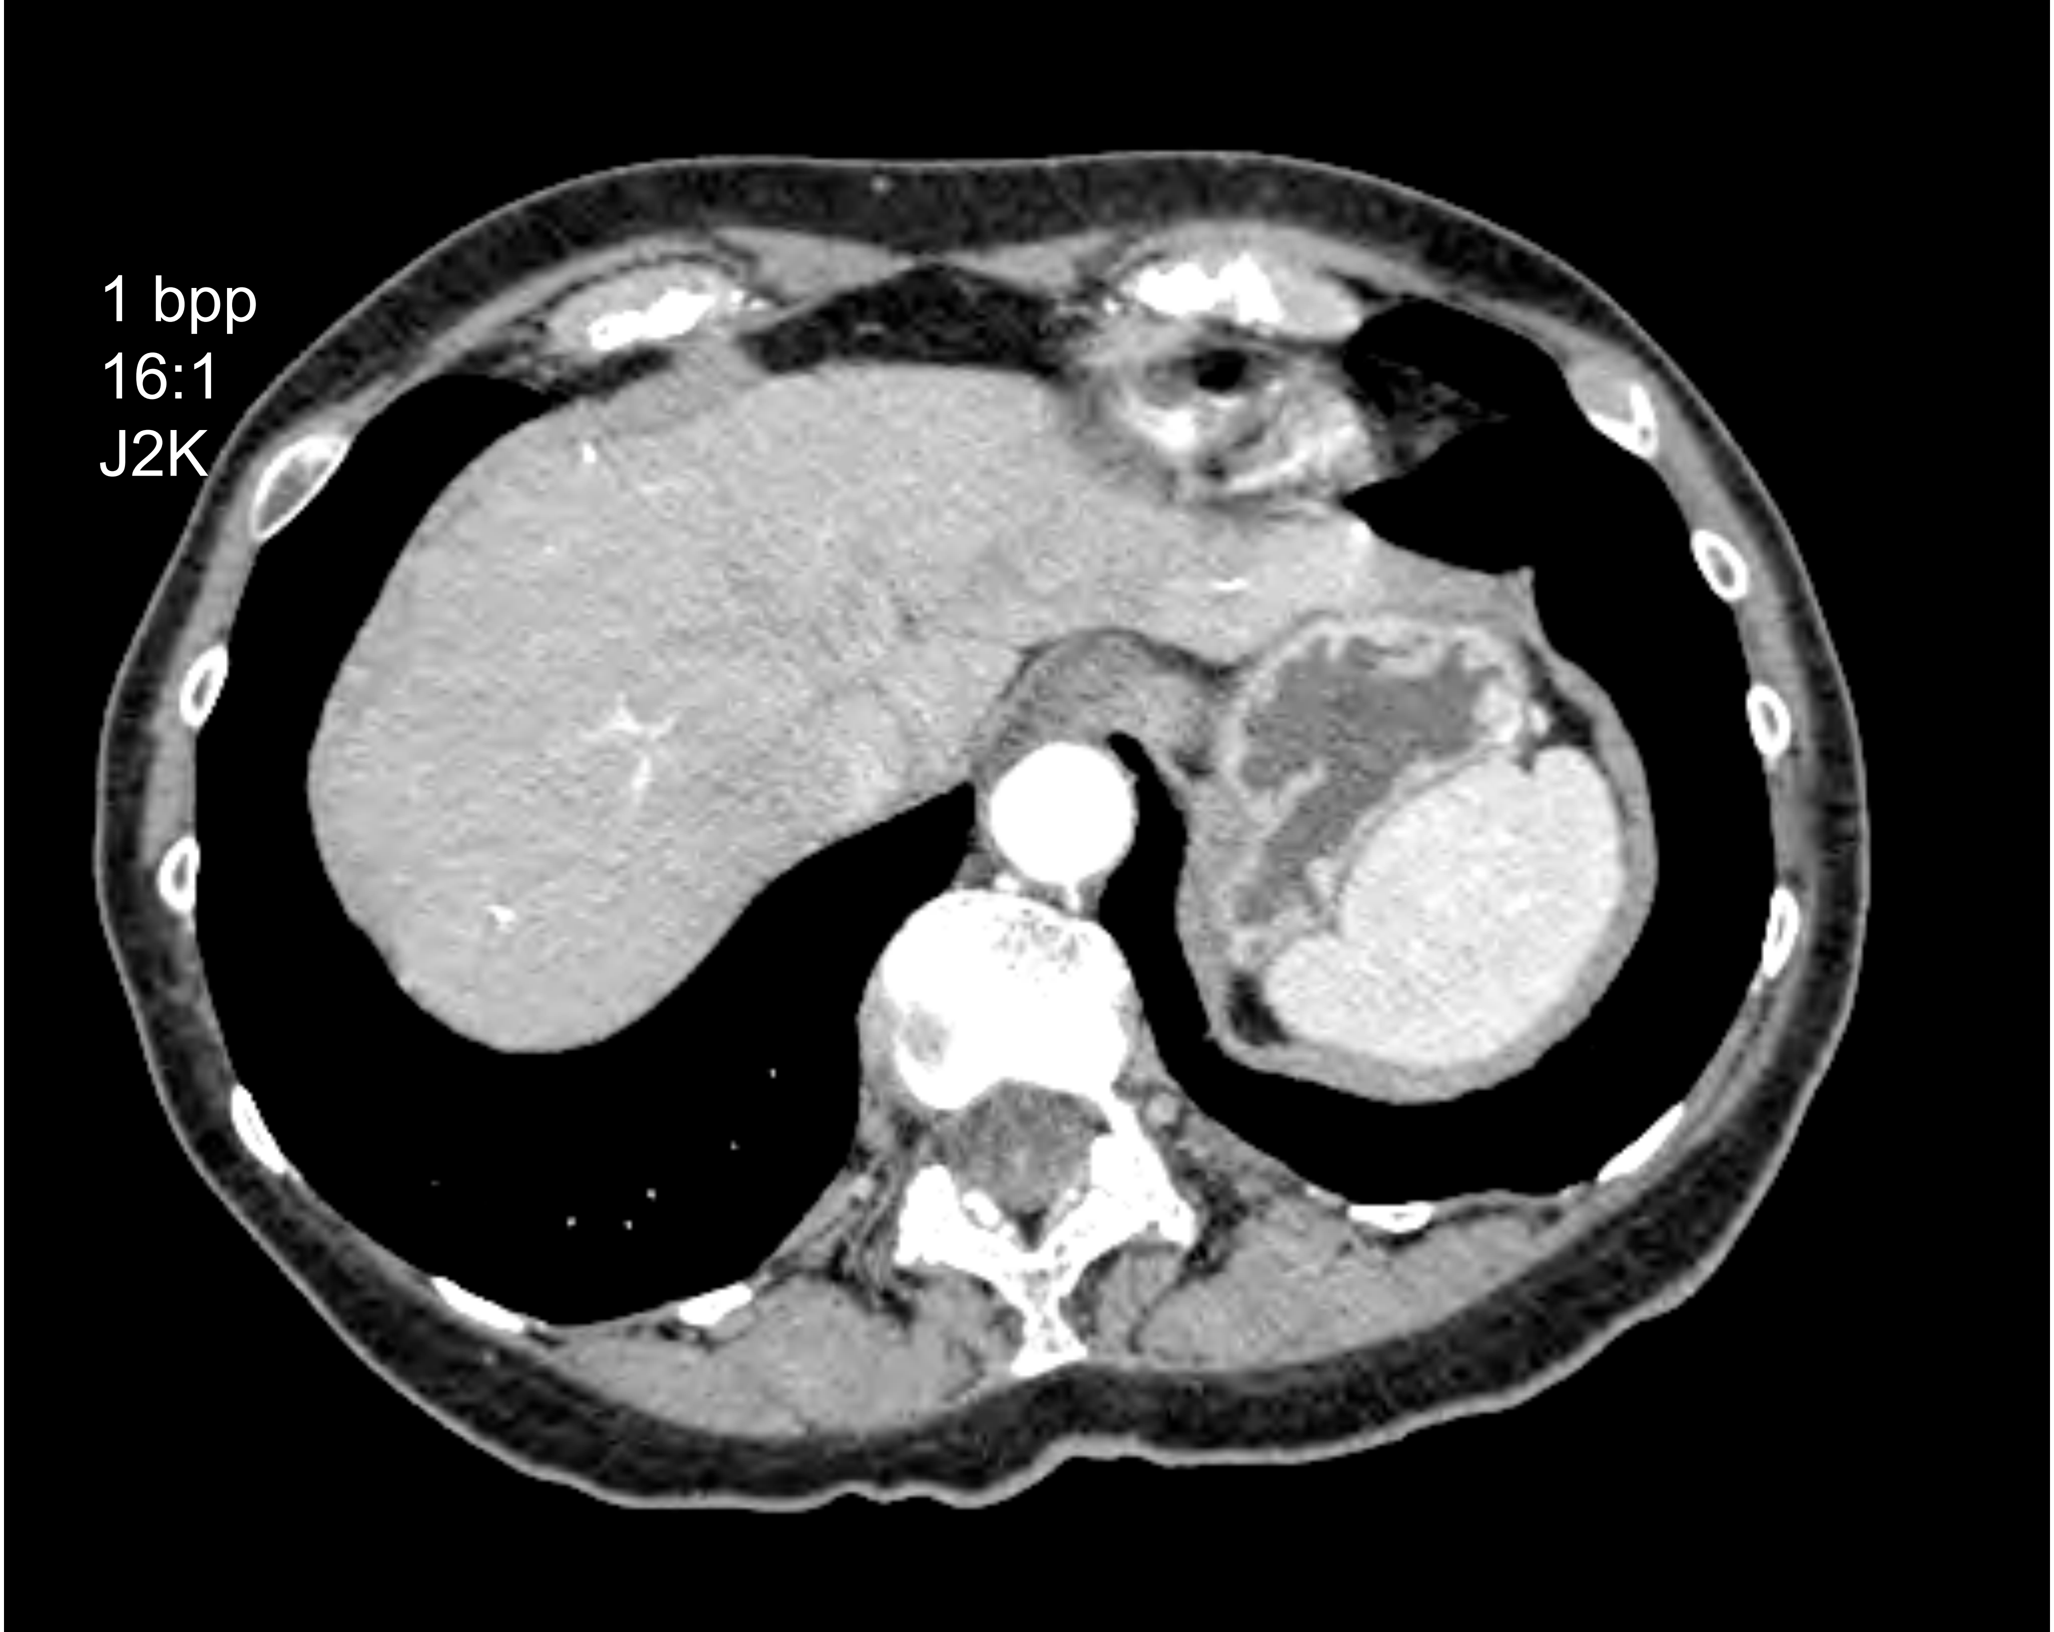
1 bpp
16:1
J2K

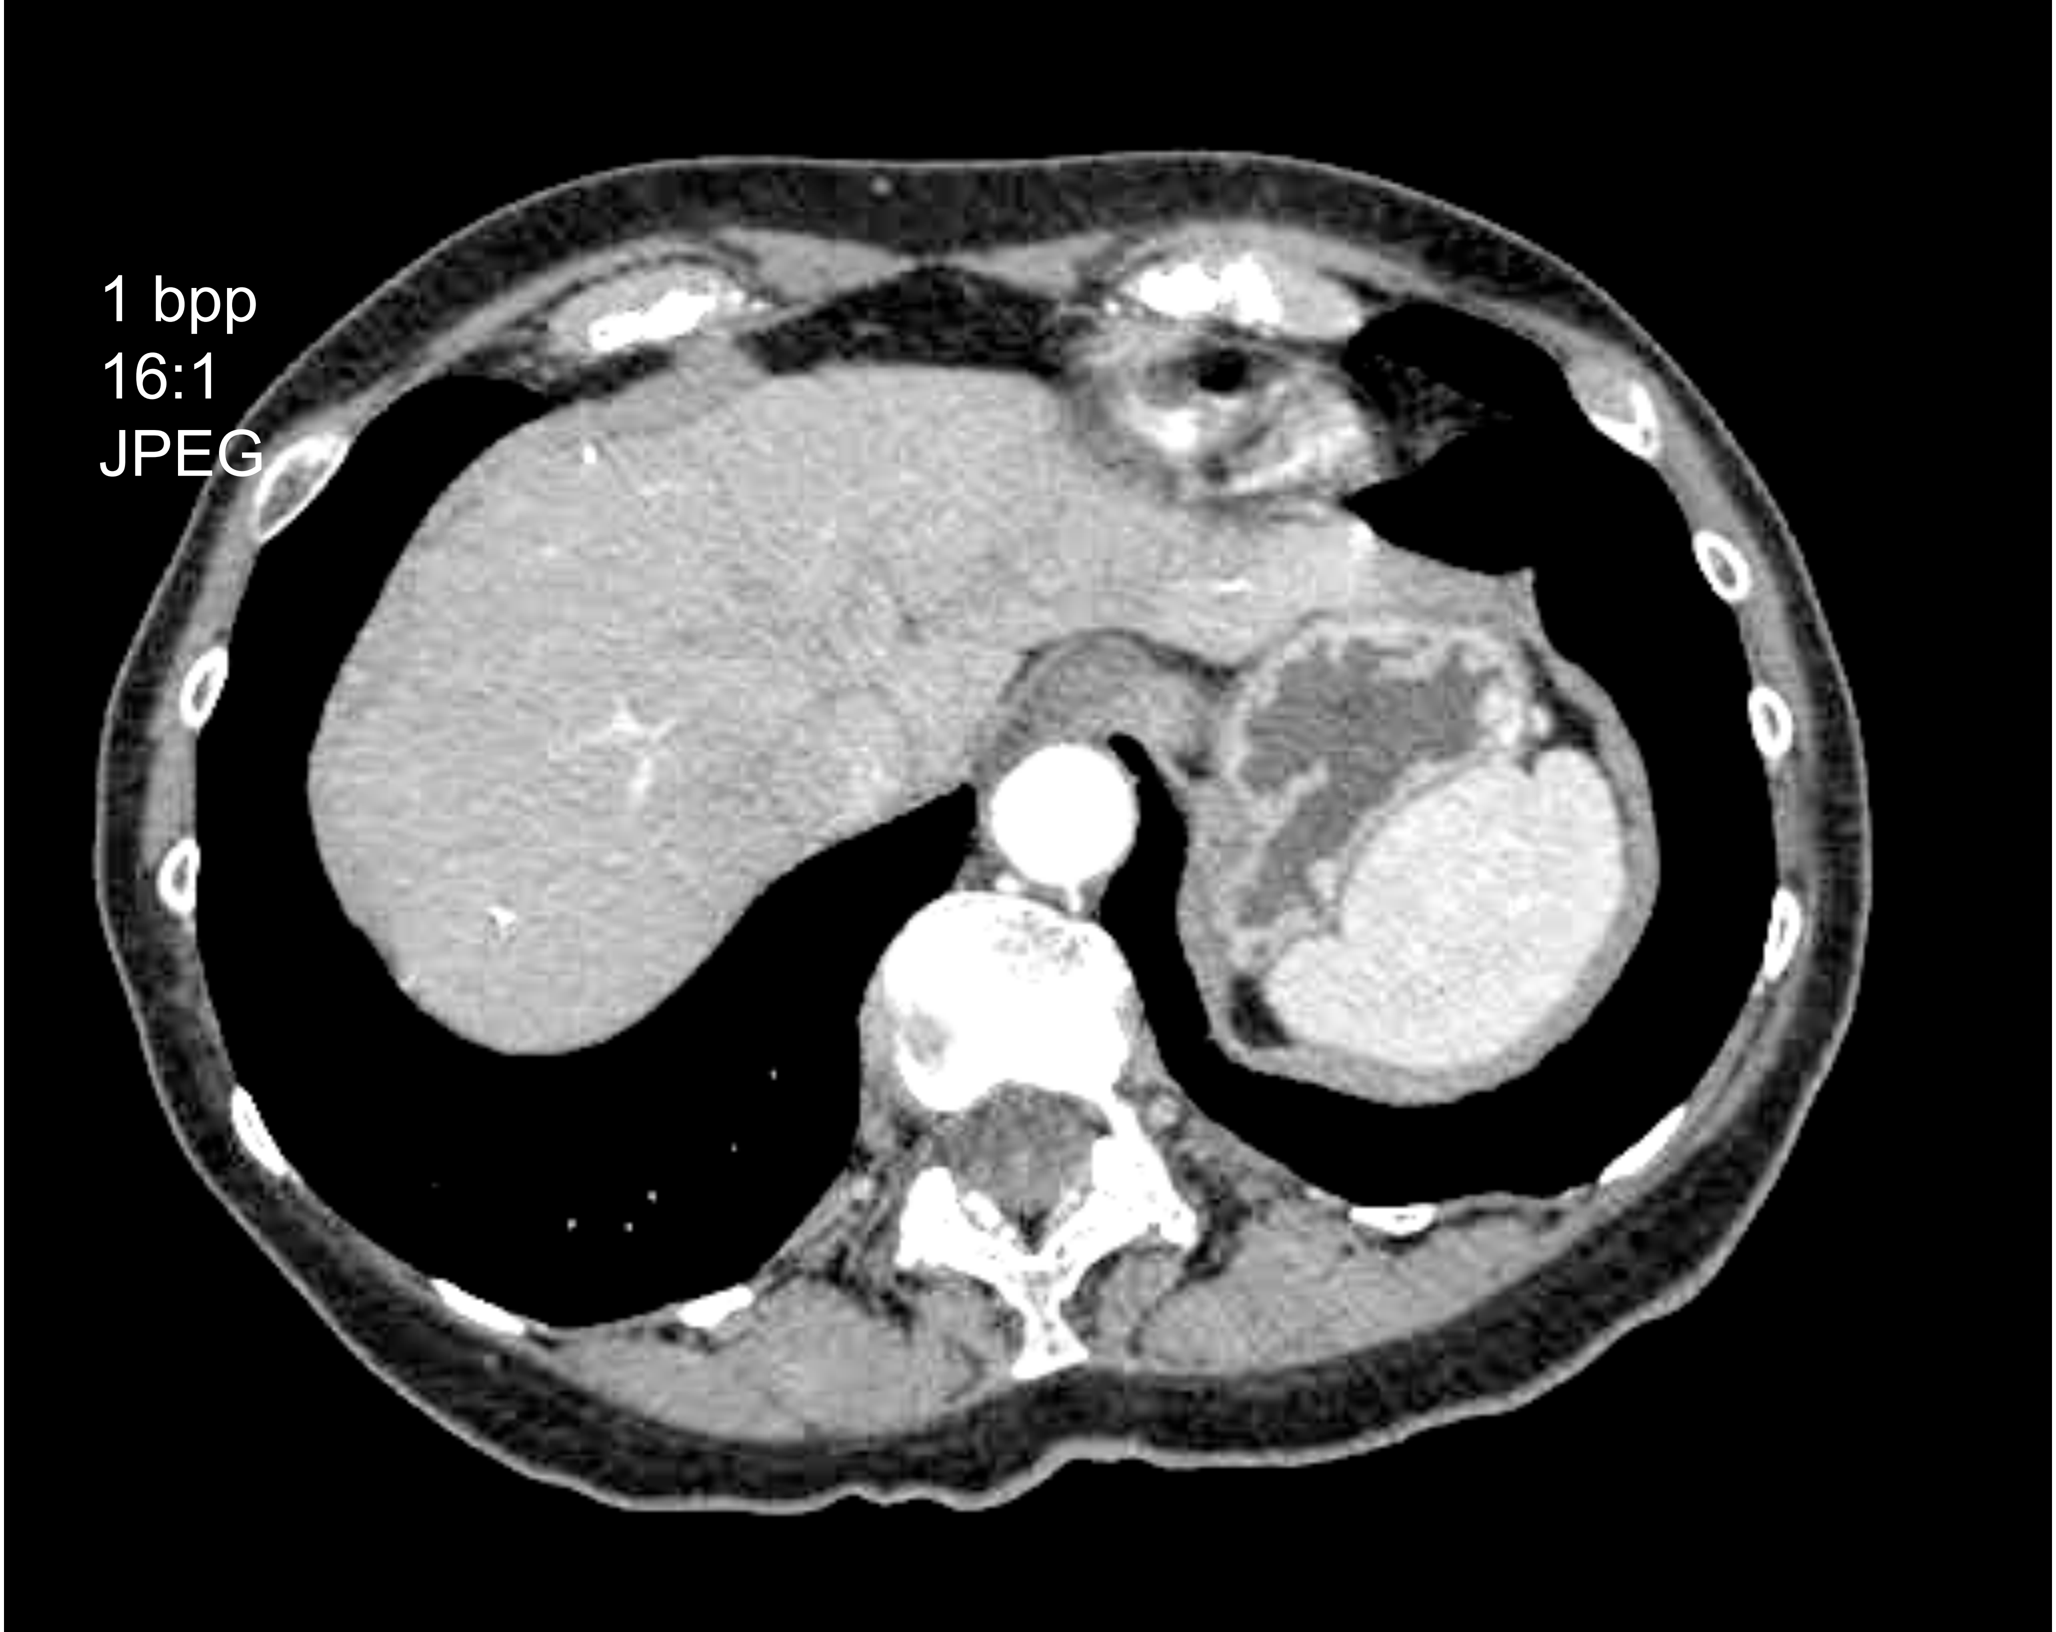
1 bpp
16:1
JPEG


0.5 bpp
32:1
J2K


0.1 bpp
160:1
J2K

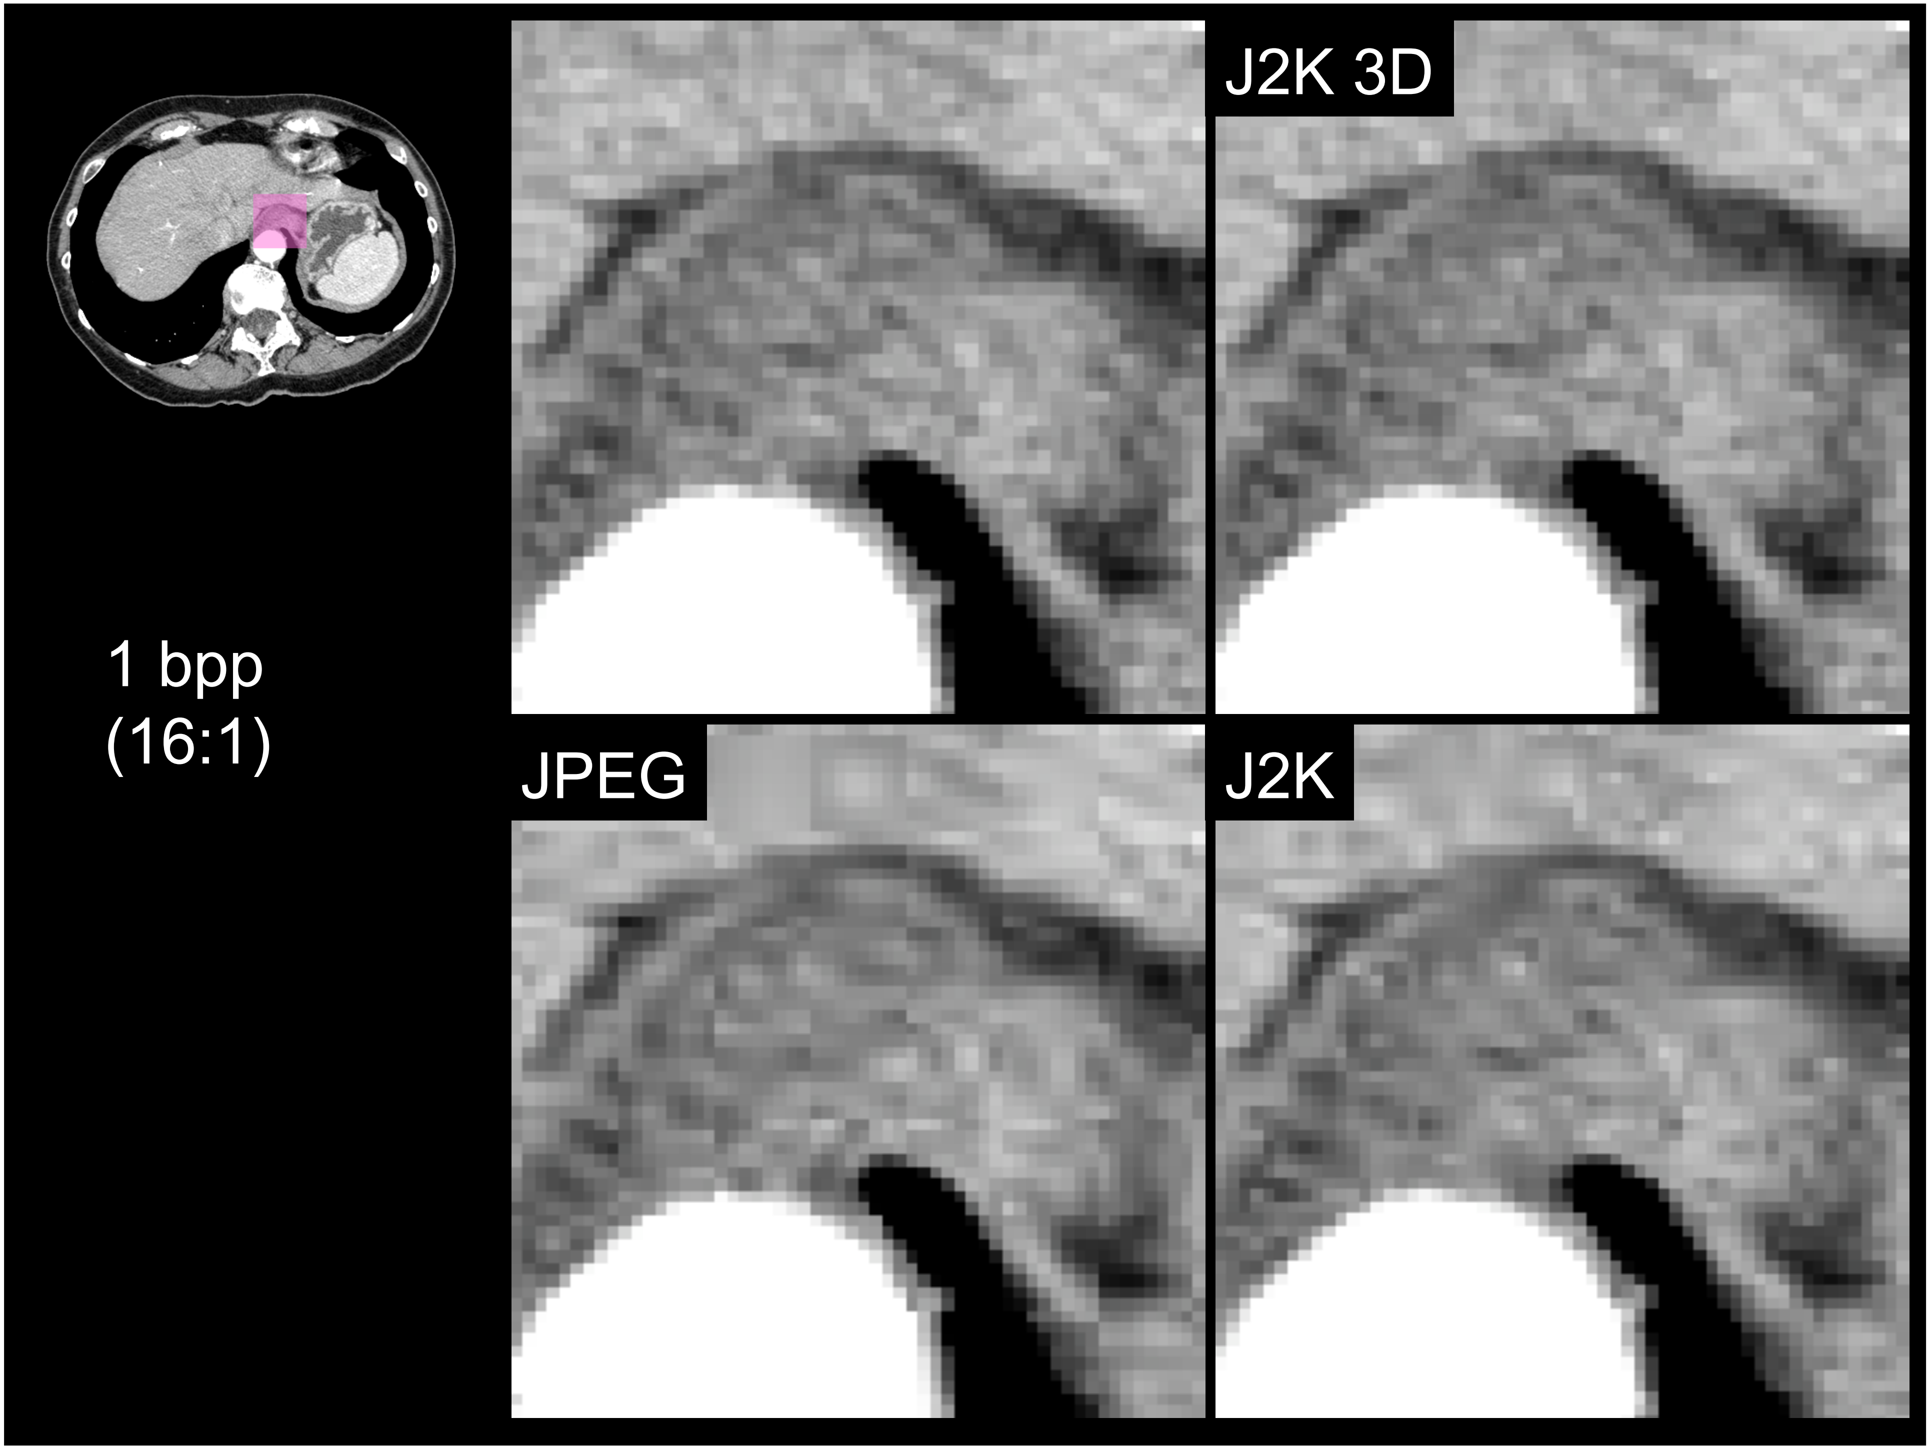
1 bpp
(16:1)
J2K 3D
JPEG
J2K

Multi-frame compression performance reality check

- Lossless compression in 3D
  - Slight gain - 15 to 20% smaller than 2D
- Lossy compression in 3D
  - Modest gain - possibly 50% smaller than 2D
  - But - only relatively modest loss before noticeable
  - Perhaps (?) 16:1
- Recent studies of JPEG 2000 on CT, 2D and 3D
  - Looked at perceived image quality & detectable difference
  - Not observer performance studies
- Need more experiments
  - Effect on observer performance unknown

Defining volumes to compress

- What to compress in 3D ?
  - Entire “volume” ?
  - Sub-sets of adjacent contiguous slices ?
- How do you find a “volume” ?
  - In a bunch of separate single frame images ?
- What is a “volume” anyway ?
  - One traversal through space
- What about other dimensions ?
  - Time (e.g. contrast phase), cardiac cycle, diffusion B value, etc. ?
- Not so easy to define a compressible volume !

DICOM & compressed volumes

- Existing DICOM CT and MR objects in common use are single frame
  - *CANNOT be used to transmit a 3D compressed volume !*
- New “Enhanced” objects are multi-frame
  - Can be used to transmit or store a 3D compressed volume
  - Presupposing frames are ordered “appropriately” (e.g., sorted by spatial location)

3D versus “multi-component”

- JPEG 2000 multi-component transform
  - Is not really “3D” per se
  - Is simply “another” dimension in which a wavelet transform can be applied
- ITU-T Rec.T.800 | ISO/IEC 15444-1 Annex J
  - “The most common multiple component transformation application is the compression of colour images … are transformed into a colour space that is more conducive to spatial compression … technique can be extended for images that have more components; for example, LANDSAT images have seven components, six of which are highly correlated … can be used for the compression of CMYK images, multiple component medical images, and any other multiple component data.”

Multi-component types

- Anything correlated between frames
- Spatial dimension
  - a single 3D volume
- Time dimension
  - contrast perfusion study
  - cardiac gated (prospectively or retrospectively)
- Other dimensions
  - Diffusion B value
  - Functional MR paradigm

True “3D” JPEG 2000

- Part 2 Annex J MCT is not the final word
- Work in progress on Part 10
  - “Extensions for three-dimensional data”
  - For “logically rectangular 3-dimensional data sets with no time component”
  - Extends MCT to support 3D “context models”
  - Goal is “moderate” improvement
- Status
  - Currently out for ballot - ends 2007/11/28
  - Core experiment report - due 2008/03/28
- Informal results - may be additional 5% improvement

Future Compression Schemes

- Need to be standards like JPEG or JPEG 2000
  - DICOM will not adopt proprietary schemes
- JPEG 2000 has been disappointing
  - Complex, slow, little consumer industry support
  - Not in browsers, not in digital cameras
- JPEG XR
  - Microsoft HD Photo (Windows Media Photo)
  - Supposedly better than JPEG, faster than JPEG 2000
  - Supports > 8 bit grayscale images (signed, floats)
  - Like JPEG is blocked based, but overlaps block edges (Lapped Bi-orthogonal Transform)

Conclusion

- Modern lossless schemes perform well
- Lossy compression remains unproven
- Plausible that a little lossy compression won't do too much harm (esp. large matrix images)
- JPEG 2000 not necessarily better than JPEG
- 3D JPEG 2000 offers modest improvement
- Need better (observer performance) studies
- Need valid use-cases to justify risk